The syslog-ng Store Box 4 F1 Administrator Guide

Publication date May 20, 2015

Abstract
This document is the primary manual of the syslog-ng Store Box 4 F1.

Copyright © 1996-2015 BalaBit S.a.r.l.

This guide is published under the Creative Commons Attribution-Noncommercial-No Derivative Works (by-nc-nd) 3.0 license. See *Appendix H, Creative Commons Attribution Non-commercial No Derivatives (by-nc-nd) License (p. 254)* for details. The latest version is always available at *https://www.balabit.com/support/documentation*.

This product includes software developed by the OpenSSL Project for use in the OpenSSL Toolkit (*http://www.openssl.org/*). This product includes cryptographic software written by Eric Young (eay@cryptsoft.com)

This documentation and the product it describes are considered protected by copyright according to the applicable laws.

The Zorp™ name and the Zorp™ logo are registered trademarks of BalaBit.

The BalaBit Shell Control Box™ name and the BalaBit Shell Control Box™ logo are registered trademarks of BalaBit.

The syslog-ng™ name and the syslog-ng™ logo are registered trademarks of BalaBit.

The BalaBit™ name and the BalaBit™ logo are registered trademarks of BalaBit.

Linux™ is a registered trademark of Linus Torvalds.

Debian™ is a registered trademark of Software in the Public Interest Inc.

Windows™ 95, 98, ME, 2000, XP, Server 2003, Vista, Server 2008 and 7 are registered trademarks of Microsoft Corporation.

MySQL™ is a registered trademark of MySQL AB in the United States, the European Union and other countries.

Oracle™, JD Edwards™, PeopleSoft™, and Siebel™ are registered trademarks of Oracle Corporation and/or its affiliates.

VMware™, VMware ESX™ and VMware View™ are trademarks or registered trademarks of VMware, Inc. and/or its affiliates.

Citrix®, ICA® and XenApp™ are trademarks or registered trademarks of Citrix Systems, Inc.

All other product names mentioned herein are the trademarks of their respective owners.

Some rights reserved.

DISCLAIMER

BalaBit is not responsible for any third-party Web sites mentioned in this document. BalaBit does not endorse and is not responsible or liable for any content, advertising, products, or other material on or available from such sites or resources. BalaBit will not be responsible or liable for any damage or loss caused or alleged to be caused by or in connection with use of or reliance on any such content, goods, or services that are available on or through any such sites or resources.

Table of Contents

Preface ..... **xi**
1. Summary of contents xi
2. Target audience and prerequisites xii
3. Products covered in this guide xii
4. Typographical conventions xiii
5. Contact and support information xiii
5.1. Sales contact xiv
5.2. Support contact xiv
5.3. Training xiv
6. About this document xiv
6.1. Summary of changes xiv
6.2. Feedback xvii
1. Introduction **1**
1.1. What SSB is 1
1.2. What SSB is not 1
1.3. Why is SSB needed 2
1.4. Who uses SSB 2
2. The concepts of SSB **3**
2.1. The philosophy of SSB 3
2.2. Collecting logs with SSB 4
2.3. Managing incoming and outgoing messages with flow-control 5
2.3.1. Flow-control and multiple destinations 7
2.4. Receiving logs from a secure channel 7
2.5. Reliable Log Transfer Protocol™ 8
2.6. Network interfaces 8
2.7. High Availability support in SSB 9
2.8. Firmware in SSB 9
2.8.1. Firmwares and high availability 9
2.9. Versions and releases of SSB 9
2.10. Licenses 10
2.11. The structure of a log message 10
2.11.1. BSD-syslog or legacy-syslog messages 10
2.11.2. IETF-syslog messages 12
3. The Welcome Wizard and the first login **16**
3.1. The initial connection to SSB 16
3.1.1. Creating an alias IP address (Microsoft Windows) 17
3.1.2. Creating an alias IP address (Linux) 20
3.1.3. Modifying the IP address of SSB 21
3.2. Configuring SSB with the Welcome Wizard 22
4. Basic settings **31**
4.1. Supported web browsers and operating systems 31
4.2. The structure of the web interface 31
4.2.1. Elements of the main workspace 34
4.2.2. Multiple web users and locking 35
4.2.3. Web interface timeout 35

4.3. Network settings 36
4.3.1. Configuring the management interface 37
4.3.2. Routing management traffic to the management interface 39
4.4. Date and time configuration 40
4.4.1. Configuring a time (NTP) server 41
4.5. SNMP and e-mail alerts 41
4.5.1. Configuring e-mail alerts 42
4.5.2. Configuring SNMP alerts 43
4.5.3. Querying SSB status information using agents 45
4.6. Configuring system monitoring on SSB 47
4.6.1. Configuring monitoring 48
4.6.2. Health monitoring 49
4.6.3. Preventing disk space fill up 49
4.6.4. Configuring message rate alerting 50
4.6.5. System related traps 52
4.6.6. Alerts related to syslog-ng 54
4.7. Data and configuration backups 55
4.7.1. Creating a backup policy 55
4.7.2. Creating configuration backups 57
4.7.3. Creating data backups 58
4.7.4. Encrypting configuration backups with GPG 58
4.8. Archiving and cleanup 59
4.8.1. Creating an archive policy 59
4.8.2. Archiving the collected data 62
4.9. Parameters of the backup and archiving protocols 62
4.9.1. Configuring Rsync over SSH 62
4.9.2. Configuring SMB 64
4.9.3. Configuring NFS 66
4.9.4. Ownership of the backup files 67
5. User management and access control 68
5.1. Managing SSB users locally 68
5.2. Setting password policies for local users 69
5.3. Managing local usergroups 71
5.4. Managing SSB users from an LDAP database 72
5.5. Authenticating users to a RADIUS server 76
5.6. Managing user rights and usergroups 77
5.6.1. Modifying group privileges 78
5.6.2. Creating new usergroups for the SSB web interface 79
5.6.3. Finding specific usergroups 80
5.6.4. How to use usergroups 80
5.6.5. Built-in usergroups of SSB 81
5.7. Listing and searching configuration changes 82
6. Managing SSB 84
6.1. Controlling SSB — restart, shutdown 84
6.2. Managing a high availability SSB cluster 85
6.2.1. Adjusting the synchronization speed 89
6.2.2. Asynchronous data replication 89
6.2.3. Redundant heartbeat interfaces 90

6.2.4. Next-hop router monitoring 92
6.3. Upgrading SSB 94
6.3.1. Upgrade checklist 94
6.3.2. Upgrading SSB (single node) 95
6.3.3. Upgrading an SSB cluster 97
6.3.4. Troubleshooting 98
6.3.5. Reverting to an older firmware version 98
6.3.6. Updating the SSB license 99
6.3.7. Exporting the configuration of SSB 100
6.3.8. Importing the configuration of SSB 102
6.4. Accessing the SSB console 103
6.4.1. Using the console menu of SSB 103
6.4.2. Enabling SSH access to the SSB host 104
6.4.3. Changing the root password of SSB 106
6.5. Sealed mode 107
6.5.1. Disabling sealed mode 107
6.6. Out-of-band management of SSB 107
6.6.1. Configuring the IPMI interface 108
6.7. Managing the certificates used on SSB 110
6.7.1. Generating certificates for SSB 112
6.7.2. Uploading external certificates to SSB 113
6.7.3. Generating TSA certificate with Windows Certificate Authority 115
6.8. Creating hostlist policies 122
6.8.1. Creating hostlists 122
6.8.2. Importing hostlists from files 123
7. Configuring message sources 126
7.1. Default message sources in SSB 126
7.2. Receiving SNMP messages 127
7.3. Creating syslog message sources in SSB 128
7.4. Creating SQL message sources in SSB 131
7.4.1. Fetching the SQL database 132
7.4.2. Configuring message parts in Basic mode 133
7.4.3. Configuring message parts in Advanced mode 136
7.4.4. Creating a fetch query manually 138
8. Storing messages on SSB 140
8.1. Default logspaces in SSB 140
8.2. Configuring the indexer 140
8.2.1. Limitations of the indexer 141
8.3. Using logstores 141
8.3.1. Viewing encrypted logs with logcat 142
8.4. Creating custom message spaces in SSB 143
8.4.1. Creating a new logstore 143
8.4.2. Creating a new text logspace 146
8.5. Managing log spaces 149
8.6. Accessing log files across the network 150
8.6.1. Sharing log files in standalone mode 150
8.6.2. Sharing log files in domain mode 152
8.6.3. Accessing shared files 155

9. Forwarding messages from SSB 157
9.1. Forwarding log messages to SQL databases 157
9.2. SQL templates in SSB 160
9.2.1. The Legacy template 160
9.2.2. The Full template 160
9.2.3. The Custom template 161
9.3. Forwarding log messages to remote servers 161
9.4. Forwarding log messages to SNMP destinations 164
10. Managing log paths 166
10.1. Default logpaths in SSB 166
10.2. Creating new log paths 166
10.3. Filtering messages 169
10.3.1. Modifying messages using rewrite 170
11. Configuring syslog-ng options 173
11.1. General syslog-ng settings 173
11.2. Timestamping configuration on SSB 175
11.3. Using name resolution on SSB 176
11.4. Setting the certificates used in TLS-encrypted log transport 177
12. Browsing log messages 181
12.1. Using the search interface 181
12.1.1. Customizing columns of the log message search interface 184
12.1.2. Metadata collected about log messages 185
12.1.3. Using wildcards and boolean search 187
12.2. Browsing encrypted log spaces 190
12.2.1. Using persistent decryption keys 191
12.2.2. Using session-only decryption keys 193
12.2.3. Assigning decryption keys to a logstore 194
12.3. Creating custom statistics from log data 195
12.3.1. Displaying log statistics 195
12.3.2. Creating reports from custom statistics 196
13. Browsing the internal messages of SSB 198
13.1. Using the internal search interface 199
13.1.1. Filtering 200
13.1.2. Exporting the results 200
13.1.3. Customizing columns of the internal search interface 200
13.2. Changelogs of SSB 201
13.3. Configuration changes of syslog-ng peers 202
13.4. Log message alerts 203
13.5. Notifications on archiving and backups 204
13.6. Statistics collection options 204
13.7. Reports 205
13.7.1. Contents of the default reports 207
13.7.2. Generating partial reports 207
13.7.3. Configuring custom reports 208
14. Classifying messages with pattern databases 211
14.1. The structure of the pattern database 212
14.2. How pattern matching works 213
14.3. Searching for rulesets 213

14.4. Creating new rulesets and rules 214
14.5. Exporting databases and rulesets 216
14.6. Importing pattern databases 216
14.7. Using pattern parsers 217
14.8. Using parser results in filters and templates 218
14.9. Using the values of pattern parsers in filters and templates 220
15. The SSB RPC API 221
15.1. Requirements for using the RPC API 221
15.2. RPC client requirements 221
15.3. Documentation of the RPC API 221
16. Troubleshooting SSB 222
16.1. Network troubleshooting 222
16.2. Gathering data about system problems 223
16.3. Viewing logs on SSB 223
16.4. Collecting logs and system information for error reporting 224
16.5. Status history and statistics 226
16.5.1. Displaying custom syslog-ng statistics 227
16.6. Troubleshooting an SSB cluster 228
16.6.1. Understanding SSB cluster statuses 228
16.6.2. Recovering SSB if both nodes broke down 230
16.6.3. Recovering from a split brain situation 231
16.6.4. Replacing a node in an SSB HA cluster 233
16.6.5. Resolving an IP conflict between cluster nodes 234
16.7. Restoring SSB configuration and data 235
Appendix A. Package contents inventory 237
Appendix B. syslog-ng Store Box Hardware Installation Guide 238
B.1. Installing the SSB hardware 238
B.2. Installing two SSB units in HA mode 240
Appendix C. Hardware specifications 241
Appendix D. syslog-ng Store Box Software Installation Guide 242
D.1. Installing the SSB software 242
Appendix E. syslog-ng Store Box VMware Installation Guide 244
E.1. Limitations of SSB under VMware 244
E.2. Installing SSB under VMware ESXi/ESX 244
E.3. Modifying the virtual disk size under VMware 245
Appendix F. syslog-ng Store Box Hyper-V Installation Guide 246
F.1. Limitations of SSB under Hyper-V 246
F.2. Installing SSB under Hyper-V 246
Appendix G. License contract for BalaBit Product 248
Appendix H. Creative Commons Attribution Non-commercial No Derivatives (by-nc-nd) License 254
Glossary 259
Index 263
List of SSB web interface labels 277

List of Examples

4.1. Number of hosts and senders 34
4.2. Creating an early time alert 51
4.3. Using the master alert to indicate unexpected events 51
4.4. Configuring NFS on the remote server 67
7.1. SQL source fetch_query 139
7.2. Query to fetch the last UID from the table 139
8.1. Mounting a shared log space using NFS on Linux 156
12.1. Searching for exact matches 187
12.2. Searching specific parts of messages 188
12.3. Combining keywords in search 188
12.4. Using parentheses in search 188
12.5. Using wildcard ? in search 188
12.6. Using wildcard * in search 189
12.7. Using combined wildcards in search 190
12.8. Searching for special characters 190
14.1. Pattern parser syntax 217
14.2. Using the STRING and ESTRING parsers 218
14.3. Patterns for multiline messages 218

List of Procedures

2.2. Collecting logs with SSB 4
3.1.1. Creating an alias IP address (Microsoft Windows) 17
3.1.2. Creating an alias IP address (Linux) 20
3.1.3. Modifying the IP address of SSB 21
3.2. Configuring SSB with the Welcome Wizard 22
4.3.1. Configuring the management interface 37
4.3.2. Routing management traffic to the management interface 39
4.4.1. Configuring a time (NTP) server 41
4.5.1. Configuring e-mail alerts 42
4.5.2. Configuring SNMP alerts 43
4.5.3. Querying SSB status information using agents 45
4.6.1. Configuring monitoring 48
4.6.3. Preventing disk space fill up 49
4.6.4. Configuring message rate alerting 50
4.7.1. Creating a backup policy 55
4.7.2. Creating configuration backups 57
4.7.3. Creating data backups 58
4.7.4. Encrypting configuration backups with GPG 58
4.8.1. Creating an archive policy 59
4.8.2. Archiving the collected data 62
4.9.1. Configuring Rsync over SSH 62
4.9.2. Configuring SMB 64
4.9.3. Configuring NFS 66
5.1. Managing SSB users locally 68
5.2. Setting password policies for local users 69
5.3. Managing local usergroups 71
5.4. Managing SSB users from an LDAP database 72
5.5. Authenticating users to a RADIUS server 76
5.6.1. Modifying group privileges 78
5.6.2. Creating new usergroups for the SSB web interface 79
6.2.3. Redundant heartbeat interfaces 90
6.2.4. Next-hop router monitoring 92
6.3.2. Upgrading SSB (single node) 95
6.3.3. Upgrading an SSB cluster 97
6.3.5. Reverting to an older firmware version 98
6.3.6. Updating the SSB license 99
6.3.7. Exporting the configuration of SSB 100
6.3.8. Importing the configuration of SSB 102
6.4.2. Enabling SSH access to the SSB host 104
6.4.3. Changing the root password of SSB 106
6.5.1. Disabling sealed mode 107
6.6.1. Configuring the IPMI interface 108
6.7.1. Generating certificates for SSB 112
6.7.2. Uploading external certificates to SSB 113
6.7.3. Generating TSA certificate with Windows Certificate Authority 115

6.8.1. Creating hostlists 122
6.8.2. Importing hostlists from files 123
7.2. Receiving SNMP messages 127
7.3. Creating syslog message sources in SSB 128
7.4.1. Fetching the SQL database 132
7.4.2. Configuring message parts in Basic mode 133
7.4.3. Configuring message parts in Advanced mode 136
8.4.1. Creating a new logstore 143
8.4.2. Creating a new text logspace 146
8.6.1. Sharing log files in standalone mode 150
8.6.2. Sharing log files in domain mode 152
9.1. Forwarding log messages to SQL databases 157
9.3. Forwarding log messages to remote servers 161
9.4. Forwarding log messages to SNMP destinations 164
10.2. Creating new log paths 166
10.3.1. Modifying messages using rewrite 170
11.4. Setting the certificates used in TLS-encrypted log transport 177
12.1.1. Customizing columns of the log message search interface 184
12.2.1. Using persistent decryption keys 191
12.2.2. Using session-only decryption keys 193
12.2.3. Assigning decryption keys to a logstore 194
12.3.2. Creating reports from custom statistics 196
13.1.3. Customizing columns of the internal search interface 200
13.7.2. Generating partial reports 207
13.7.3. Configuring custom reports 208
14.4. Creating new rulesets and rules 214
14.8. Using parser results in filters and templates 218
16.1. Network troubleshooting 222
16.3. Viewing logs on SSB 223
16.4. Collecting logs and system information for error reporting 224
16.5.1. Displaying custom syslog-ng statistics 227
16.6.2. Recovering SSB if both nodes broke down 230
16.6.3. Recovering from a split brain situation 231
16.6.4. Replacing a node in an SSB HA cluster 233
16.6.5. Resolving an IP conflict between cluster nodes 234
16.7. Restoring SSB configuration and data 235
B.1. Installing the SSB hardware 238
B.2. Installing two SSB units in HA mode 240
D.1. Installing the SSB software 242
E.2. Installing SSB under VMware ESXi/ESX 244
E.3. Modifying the virtual disk size under VMware 245
F.2. Installing SSB under Hyper-V 246

Preface

Welcome to the syslog-ng Store Box 4 F1 Administrator Guide!

This document describes how to configure and manage the syslog-ng Store Box (SSB). Background information for the technology and concepts used by the product is also discussed.

1. Summary of contents

Chapter 1, Introduction (p. 1) describes the main functionality and purpose of the syslog-ng Store Box.

Chapter 2, The concepts of SSB (p. 3) discusses the technical concepts and philosophies behind SSB.

Chapter 3, The Welcome Wizard and the first login (p. 16) describes what to do after assembling SSB — it is a step-by-step guide for the initial configuration.

Chapter 4, Basic settings (p. 31) provides detailed description on configuring and managing SSB as a host.

Chapter 5, User management and access control (p. 68) describes how to manage user accounts and privileges.

Chapter 6, Managing SSB (p. 84) explains the basic management tasks of SSB, including the basic control (for example, shutdown or reboot) of the appliance and upgrading.

Chapter 7, Configuring message sources (p. 126) provides description on using the built-in message sources, creating new message sources and receiving SNMP messages.

Chapter 8, Storing messages on SSB (p. 140) describes how to store log messages in log spaces.

Chapter 9, Forwarding messages from SSB (p. 157) explains how to forward log messages to remote destinations.

Chapter 10, Managing log paths (p. 166) discusses the management of log paths.

Chapter 11, Configuring syslog-ng options (p. 173) describes the configuration options of the syslog-ng server running on syslog-ng Store Box.

Chapter 12, Browsing log messages (p. 181) describes how to browse logs on SSB.

Chapter 13, Browsing the internal messages of SSB (p. 198) describes how to browse internal messages and reports of SSB.

Chapter 14, Classifying messages with pattern databases (p. 211) describes how to parse and classify messages using the pattern database.

Chapter 15, The SSB RPC API (p. 221) describes how to access and query SSB logspaces from remote applications.

Chapter 16, Troubleshooting SSB (p. 222) describes troubleshooting and maintenance procedures of syslog-ng Store Box (SSB).

Appendix A, Package contents inventory (p. 237) lists the contents of the package you receive with the syslog-ng Store Box.

Appendix B, syslog-ng Store Box Hardware Installation Guide (p. 238) describes how to set up the syslog-ng Store Box (SSB) hardware.

Appendix C, Hardware specifications (p. 241) describes the hardware specifications of the syslog-ng Store Box (SSB) appliance.

Appendix D, syslog-ng Store Box Software Installation Guide (p. 242) describes how to install syslog-ng Store Box (SSB) on certified hardware.

Appendix E, syslog-ng Store Box VMware Installation Guide (p. 244) describes how to install syslog-ng Store Box (SSB) as a VMware virtual appliance.

Appendix G, License contract for BalaBit Product (p. 248) includes the text of the End User License Agreement applicable to SSB products.

Appendix H, Creative Commons Attribution Non-commercial No Derivatives (by-nc-nd) License (p. 254) includes the text of the Creative Commons Attribution Non-commercial No Derivatives (by-nc-nd) License applicable to The syslog-ng Store Box 4 F1 Administrator Guide.

The *Glossary (p. 259)* provides definitions of important terms used in this guide.

2. Target audience and prerequisites

This guide is intended for auditors, consultants, and security experts responsible for securing, auditing, and monitoring server administration processes, especially remote server management. It is also useful for IT decision makers looking for a tool to improve the security and auditability of their servers, or to facilitate compliance to the Sarbanes-Oxley (SOX), the Health Insurance Portability and Accountability Act (HIPAA), Basel II, or the Payment Card Industry (PCI) standard.

The following skills and knowledge are necessary for a successful SSB administrator:

- At least basic system administration knowledge.
- An understanding of networks, TCP/IP protocols, and general network terminology.
- An understanding of system logging and the protocols used in remote system logging.
- Familiarity with the concepts of the syslog-ng and the syslog-ng Agent for Windows applications.
- Working knowledge of the UNIX or Linux operating system is not mandatory but highly useful.

3. Products covered in this guide

This guide describes the use of the syslog-ng Store Box version 4 F1.

Note
Users of the syslog-ng Store Box are entitled to use the syslog-ng Premium Edition application as a log collector agent for SSB. This guide does not cover the installation and configuration of syslog-ng Premium Edition, see *the syslog-ng documentation*.

4. Typographical conventions

Before you start using this guide, it is important to understand the terms and typographical conventions used in the documentation. For more information on specialized terms and abbreviations used in the documentation, see the Glossary at the end of this document.

The following kinds of text formatting and icons identify special information in the document.

Tip
Tips provide best practices and recommendations.

Note
Notes provide additional information on a topic, and emphasize important facts and considerations.

Warning
Warnings mark situations where loss of data or misconfiguration of the device is possible if the instructions are not obeyed.

`Command`	Commands you have to execute.
Emphasis	Reference items, additional readings.
`/path/to/file`	File names.
`Parameters`	Parameter and attribute names.
Label	GUI output messages or dialog labels.
Menu	A submenu or menu item in the menu bar.
Button	Buttons in dialog windows.

5. Contact and support information

This product is developed and maintained by BalaBit-Europe. We are located in Budapest, Hungary. Our address is:

BalaBit-Europe
2 Alíz Street
H-1117 Budapest, Hungary
Tel: +36 1 398-6700
Fax: +36 1 208-0875
E-mail: <info@balabit.com>

Web: *https://www.balabit.com/*

5.1. Sales contact

You can directly contact us with sales related topics at the e-mail address <sales@balabit.com>, or *leave us your contact information and we call you back*.

5.2. Support contact

To access the BalaBit Online Support System (BOSS), sign up for an account at *the MyBalaBit page* and *request access to the BalaBit Online Support System (BOSS)*. Online support is available 24 hours a day.

BOSS is available only for registered users with a valid support package.

Support e-mail address: <support@balabit.com>.

Support hotline: +36 1 398 6700 (available from 9 AM to 5 PM CET on weekdays)

5.3. Training

BalaBit-Europe holds courses on using its products for new and experienced users. For dates, details, and application forms, visit the *https://my.balabit.com/training/* webpage.

6. About this document

This guide is a work-in-progress document with new versions appearing periodically.

The latest version of this document can be downloaded from the BalaBit website *here*.

6.1. Summary of changes

6.1.1. Version 3 LTS - 3 F2

Changes in product:

- New RPC API, which allows accessing and querying the log messages stored on SSB from remote applications using a RESTful protocol over HTTPS.
- New log message search interface.
- Nested groups can be disabled when querying LDAP servers.
- Multiple append domains can be set in the Connection Policy.
- The internal timestamp handling of SSB has been changed to improve indexing and search performance.
- SSB is now officially supported on VMWare ESX 4.0 and later and ESXi 5.0 and later as well.

Changes in documentation:

- *Appendix C, Hardware specifications (p. 241)* has been added to the document.
- *Section 8.2.1, Limitations of the indexer (p. 141)* has been added to the document.
- Added a description of the Locked status to *Section 4.2, The structure of the web interface (p. 31)*.
- Updated descriptions and requirements in *Appendix E, syslog-ng Store Box VMware Installation Guide (p. 244)*.
- *Chapter 15, The SSB RPC API (p. 221)* has been added to the document.
- Added the list of ports required by the IPMI interface to *Appendix B, syslog-ng Store Box Hardware Installation Guide (p. 238)*.
- Updated *Chapter 12, Browsing log messages (p. 181)* with descriptions of the new log message search interface.

6.1.2. Version 3 LTS - 3 F1

Changes in product:

- SAN support has been removed from the document and the product.
- *Section 7.4, Creating SQL message sources in SSB (p. 131)* has been added to the document.
- *Section 6.2.2, Asynchronous data replication (p. 89)* has been added to the document.
- The Extended schema for SQL destinations has been removed from the product and from the documentation in SSB version 3.0.1.
- Added a warning about a possible problem when archiving to Windows 2008 R2 hosts using the CIFS protocol to *Procedure 4.8.1, Creating an archive policy (p. 59)*.
- Web interface timeout description has been added to *Chapter 4, Basic settings (p. 31)*.
- Added a note about redundant HA links to *Section 16.6.1, Understanding SSB cluster statuses (p. 228)*.
- Described 'All' option for Lists statistics in *Section 12.3, Creating custom statistics from log data (p. 195)*.
- Described firmware upgrade notes in *Procedure 6.3.2, Upgrading SSB (single node) (p. 95)*.
- The description of patterns for multiline messages has been added to *Section 14.7, Using pattern parsers (p. 217)*.

Changes in documentation:

- *Appendix G, License contract for BalaBit Product (p. 248)* has been updated.

6.1.3. Version 2 F1 - 3 LTS

Changes in product:

- Description of the **Don't parse messages** has been added to *Procedure 7.3, Creating syslog message sources in SSB (p. 128)*.

- Figures *Configuring SNMP and e-mail alerting*; *Default message sources in SSB*; *Configuring syslog-ng options*; *Configuring persistent name resolution*; *Configuring TLS settings for syslog-ng*; *Creating database destinations*; *Creating server destinations*; *Displaying search information*;

Displaying statistics; *Creating a new logstore*; *Creating a new text logspace*; *Creating new message sources* have been updated.

- Figures *Creating an early time alert*; *Using the master alert to indicate unexpected events* and *Modifying messages using rewrite* have been added to the document.
- Alerts *Message rate was outside the specified limits* and *Too many message rate alerts were generated* have been added to *Section 4.6.6, Alerts related to syslog-ng (p. 54)*.
- The list of sources in *Section 7.1, Default message sources in SSB (p. 126)* has been updated with the BSD-syslog (legacy TCP) protocol.
- *Procedure 4.6.4, Configuring message rate alerting (p. 50)* has been added to the document.
- *Section 12.1.3, Using wildcards and boolean search (p. 187)* has been added to the document.
- *Procedure 10.3.1, Modifying messages using rewrite (p. 170)* has been added to the document.
- *Procedure 4.6.3, Preventing disk space fill up (p. 49)* has been added to the document.
- *Procedure 16.7, Restoring SSB configuration and data (p. 235)* has been added to the document.
- *Section 7.1, Default message sources in SSB (p. 126)* has been updated with new default ports.
- *Section 4.1, Supported web browsers and operating systems (p. 31)* has been updated with new supported and tested browsers.
- *Section 13.6, Statistics collection options (p. 204)* has been updated.
- Rate limiting has been removed from *Procedure 7.3, Creating syslog message sources in SSB (p. 128)*.

Changes in documentation:

- The document chapters have been restructured. *Section 1, Summary of contents (p. xi)* has been updated with the new structure.
- Screenshots in the HTML version of the document have been resized for better visibility.
- The troubleshooting section has been moved to *Chapter 16, Troubleshooting SSB (p. 222)*.
- A warning has been added to *Procedure 4.5.1, Configuring e-mail alerts (p. 42)*.

6.1.4. Version 2 LTS - 2 F1

Changes in product:

- *Procedure 9.4, Forwarding log messages to SNMP destinations (p. 164)* has been added to the document.
- References to hardware manuals in Appendix *syslog-ng Store Box Hardware Installation Guide* have been updated and corrected.
- Appendix *syslog-ng Store Box VMware Installation Guide* has been added to the document.
- Procedure *Modifying the IP address of SSB* has been added to the document.
- Section *High Availability status explained* has been updated with 'Converted' status.
- Windows 7 has been added to the supported operating systems list in sections *Supported web browsers and operating systems*, *Supported protocols and client applications* and *Viewing session information and replaying audit trails*.
- A note has been added to *Section 8.6.3, Accessing shared files (p. 155)*.

Changes in documentation:

- Section *Summary of changes* has been added to the document.
- Procedures in the HTML version of the document appear on separate HTML pages.
- Labels of cross-references pointing to procedure steps have been corrected.
- Procedures have been restructured to facilitate easier understanding.
- Latin abbreviations have been replaced in document with their English equivalents.
- Editorial changes.

6.2. Feedback

Any feedback is greatly appreciated, especially on what else this document should cover. General comments, errors found in the text, and any suggestions about how to improve the documentation is welcome at documentation@balabit.com.

Chapter 1. Introduction

This chapter introduces the syslog-ng Store Box (SSB) in a non-technical manner, discussing how and why is it useful, and what additional benefits it offers to an existing IT infrastructure.

1.1. What SSB is

SSB is a device that collects, processes, stores, monitors, and manages log messages. It is a central logserver appliance that can receive system (syslog and eventlog) log messages and Simple Network Management Protocol (SNMP) messages from your network devices and computers, store them in a trusted and signed logstore, automatically archive and backup the messages, and also classify the messages using artificial ignorance.

The most notable features of SSB are the following:

- Secure log collection using Transport Layer Security (TLS).
- Trusted, encrypted, signed, timestamped storage.
- Ability to collect log messages from a wide range of platforms, including Linux, Unix, BSD, Sun Solaris, HP-UX, IBM AIX, IBM System i, as well as Microsoft Windows XP, Server 2003, Vista, and Server 2008.
- Forwards messages to log analyzing engines.
- Classifies messages using customizable pattern databases for real-time log monitoring, alerting, and artificial ignorance.
- High Availability (HA) support to ensure continuous log collection in business-critical environments.
- Real-time log monitoring and alerting.
- Retrieves group memberships of the administrators from a Lightweight Directory Access Protocol (LDAP) database.
- Strict, yet easily customizable access control to grant users access only to selected log messages

SSB is configured and managed from any modern web browser that supports HTTPS connections, JavaScript, and cookies. Supported browsers: Mozilla Firefox 10 and Microsoft Internet Explorer 8 and 9. Other tested browsers: Mozilla Firefox 3.6 and Google Chrome 17.

1.2. What SSB is not

SSB is not a log analyzing engine, it is able to classify individual log messages using artificial ignorance, much like the popular logcheck application of the Unix world. SSB comes with a built-in feature to store log message patterns that are considered "normal". Messages matching these patterns are produced during the legitimate use of the applications (for example sendmail, Postfix, MySQL, and so on), and are unimportant from the log monitoring perspective, while the remaining messages may contain something “interesting”. The administrators can define log patterns on the SSB interface, label matching messages (for example security event, and so on.) and request alerts if a specific pattern is encountered. For thorough log analysis, SSB can also forward the incoming log messages to external log analyzing engines.

1.3. Why is SSB needed

Log messages contain information about the events happening on the hosts. Monitoring system events is essential for security and system health monitoring reasons. A well-established log management solution offers several benefits to an organization. It ensures that computer security records are stored in sufficient detail, and provides a simple way to monitor and review these logs. Routine log reviews and continuous log analysis help to identify security incidents, policy violations, or other operational problems. Logs also often form the base of auditing and forensic analysis, product troubleshooting and support. There are also several laws, regulations and industrial standards that explicitly require the central collection, periodic review, and long-time archiving of log messages. Examples to such regulations are the Sarbanes-Oxley Act (SOX), the Basel II accord, the Health Insurance Portability and Accountability Act (HIPAA), or the Payment Card Industry Data Security Standard (PCI-DSS).

Built around the popular syslog-ng application used by thousands of organizations worldwide, the syslog-ng Store Box (SSB) brings you a powerful, easy to configure appliance to collect and store your logs. Using the features of the latest syslog-ng Premium Edition to their full power, SSB allows you to collect, process, and store log messages from a wide range of platforms and devices.

All data can be stored in encrypted, digitally signed, and optionally timestamped files, preventing any modification or manipulation, satisfying the highest security standards and policy compliance requirements.

1.4. Who uses SSB

SSB is useful for everyone who has to collect, store, and review log messages. In particular, SSB is invaluable for:

- *Central log collection and archiving*: SSB offers a simple, reliable, and convenient way of collecting log messages centrally. It is essentially a high-capacity log server with high availability support. Being able to collect logs from several different platforms makes it easy to integrate into any environment.
- *Secure log transfer and storage*: Log messages often contain sensitive information and also form the base of audit trails for several applications. Preventing eavesdropping during message transfer and unauthorized access once the messages reach the logserver is essential for security and privacy reasons.
- *Policy compliance*: Many organization must comply to regulations like the Sarbanes-Oxley Act (SOX), the Basel II accord, the Health Insurance Portability and Accountability Act (HIPAA), or the Payment Card Industry Data Security Standard (PCI-DSS). These regulations often have explicit or implicit requirements about log management, such as the central collection of log messages, the use of log analysis to prevent and detect security incidents, or guaranteeing the availability of log messages for an extended period of time — up to several years. SSB helps these organizations to comply with these regulations.
- *Automated log monitoring and log preprocessing*: Monitoring log messages is an essential part of system-health monitoring and security incident detection and prevention. SSB offers a powerful platform that can classify tens of thousands of messages real-time to detect messages that deviate from regular messages, and promptly raise alerts. Although this classification does not offer as complete inspection as a log analyzing application, SSB can process much more messages than a regular log analyzing engine, and also filter out unimportant messages to decrease the load on the log analyzing application.

Chapter 2. The concepts of SSB

This chapter discusses the technical concepts of SSB.

2.1. The philosophy of SSB

The syslog-ng Store Box (SSB) is a log server appliance that collects, stores and monitors the log messages sent by network devices, applications and computers. SSB can receive traditional syslog messages, syslog messages that comply with the new Internet Engineering Task Force (IETF) standard, eventlog messages from Microsoft Windows hosts, as well as SNMP messages.

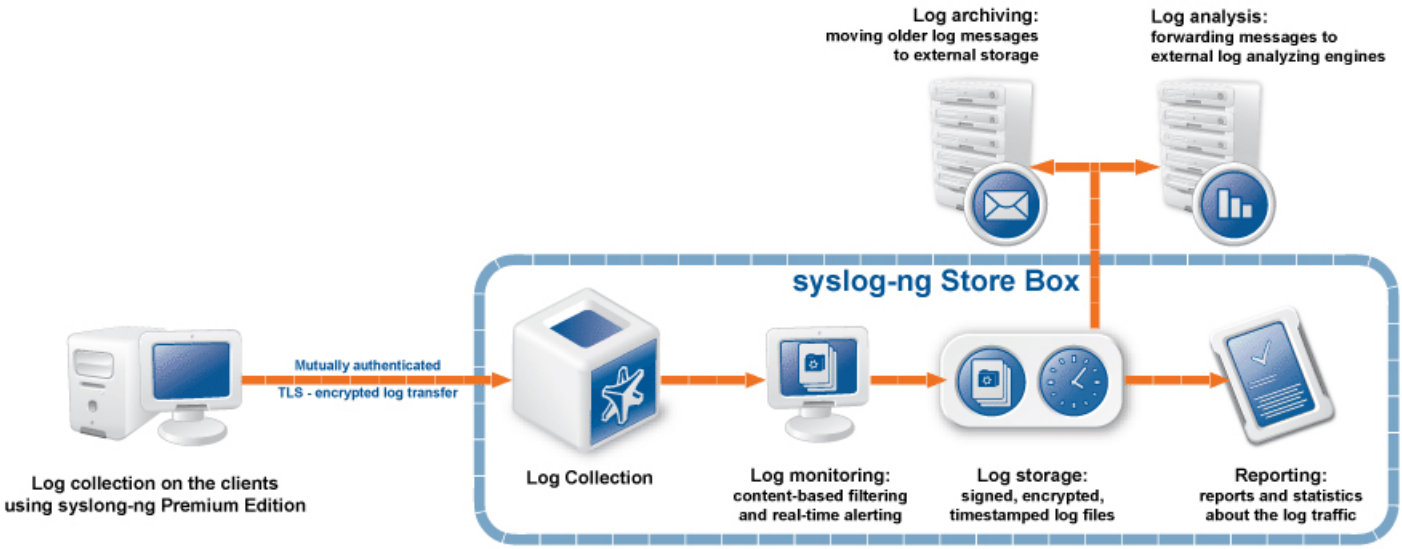

Figure 2.1. The philosophy of the syslog-ng Store Box

Clients can send messages to SSB using their own logging application if it supports the *BSD-syslog* (RFC 3164) or the *IETF-syslog* (RFC 5424-5428) protocol, or they can use the syslog-ng Premium Edition application to act as the log-forwarding agent of SSB.

The main purpose of SSB is to collect the logs from the clients and store them on its hard disk. The messages are stored in so-called logspaces. There are two types of logspaces: the first stores messages in traditional plain-text files, while the second one uses a binary format that can be compressed, encrypted, digitally signed, and also timestamped.

The syslog-ng application reads incoming messages and forwards them to the selected *destinations*. The syslog-ng application can receive messages from files, remote hosts, and other *sources*.

Log messages enter syslog-ng in one of the defined sources, and are sent to one or more *destinations*. In case of the clients, one of the destinations is the syslog-ng Store Box; the destinations on the SSB can be logspaces or remote servers, such as database servers or log analyzing engines.

Sources and destinations are independent objects; *log paths* define what syslog-ng does with a message, connecting the sources to the destinations. A log path consists of one or more sources and one or more destinations; messages arriving to a source are sent to every destination listed in the log path. A log path defined in syslog-ng is called a *log statement*.

Optionally, log paths can include *filters*. Filters are rules that select only certain messages, for example, selecting only messages sent by a specific application. If a log path includes filters, syslog-ng sends only the messages satisfying the filter rules to the destinations set in the log path.

SSB is configured by an administrator or auditor using a web browser.

2.2. Procedure – Collecting logs with SSB

Purpose:

The following procedure illustrates the route of a log message from its source on the syslog-ng client to the syslog-ng Store Box.

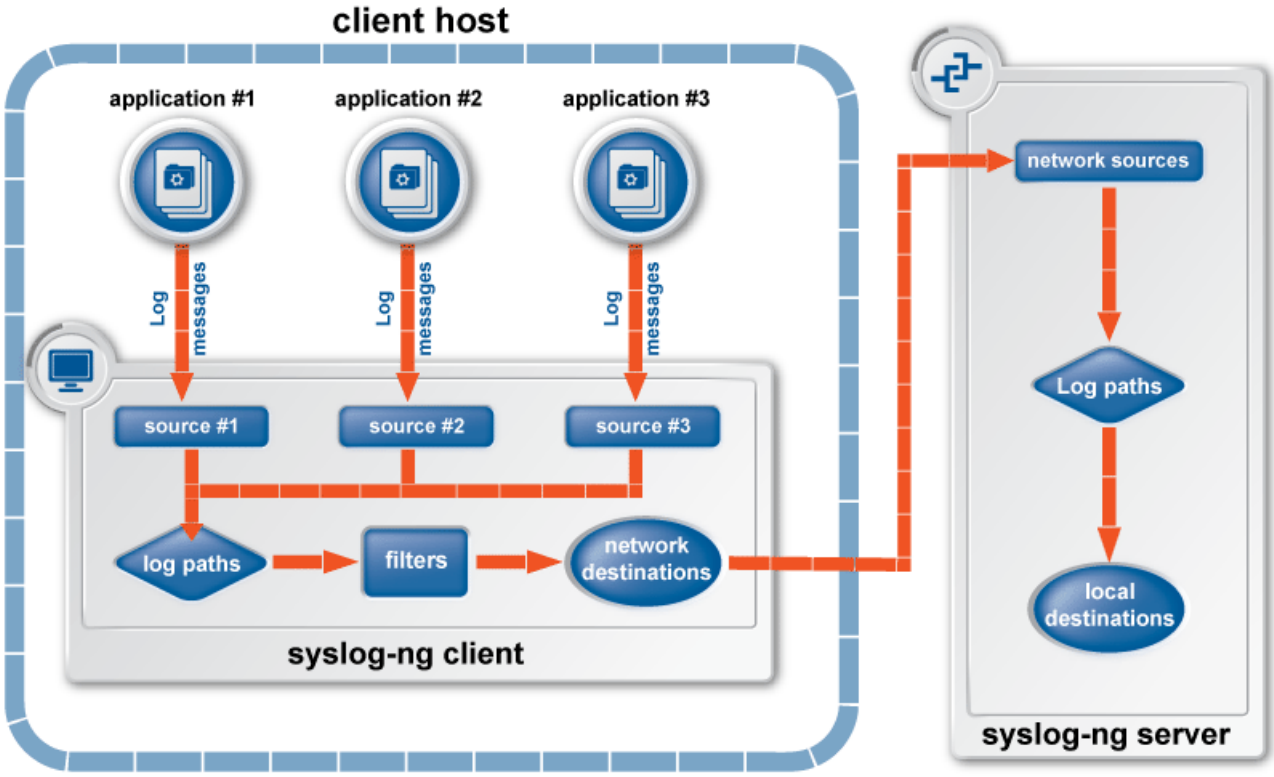

Figure 2.2. The route of a log message

Steps:

Step 1. A device or application sends a log message to a source on the syslog-ng client. For example, an Apache web server running on Linux enters a message into the `/var/log/apache` file.

Step 2. The syslog-ng client running on the web server reads the message from its `/var/log/apache` source.

Step 3. The syslog-ng client processes the first log statement that includes the `/var/log/apache` source.

Step 4. The syslog-ng client performs optional operations (for example message filtering) on the message; for example, it compares the message to the filters of the log statement (if any). If the message complies with all filter rules, syslog-ng sends the message to the destinations set in the log statement, for example, to the remote syslog-ng server.
After that, the syslog-ng client processes the next log statement that includes the `/var/log/apache` source, repeating Steps 3-4.

Step 5. The message sent by the syslog-ng client arrives to a source set on the syslog-ng Store Box.

Step 6. The syslog-ng Store Box reads the message from its source and processes the first log path that includes that source.

Step 7. The syslog-ng server performs optional operations (for example message filtering, or pattern matching to compare the message to a list of known messages). If the message complies with all filter rules, SSB sends the message to the destinations set in the log path. The destinations are local, optionally encrypted files on SSB, or remote servers such as a database server.

Step 8. SSB processes the next log statement, repeating Steps 6-8.

Note
The syslog-ng application can stop reading messages from its sources if the destinations cannot process the sent messages. This feature is called flow-control and is detailed in *Section 2.3, Managing incoming and outgoing messages with flow-control (p. 5)*.

2.3. Managing incoming and outgoing messages with flow-control

This section describes the internal message-processing model of syslog-ng, as well as the flow-control feature that can prevent message loss. To use flow-control, the **Flow** option must be enabled for the particular log path.

The syslog-ng application monitors (polls) the sources defined in its configuration file, periodically checking each source for messages. When a log message is found in one of the sources, syslog-ng polls every source and reads the available messages. These messages are processed and put into the output buffer of syslog-ng (also called fifo). From the output buffer, the operating system sends the messages to the appropriate destinations.

In large-traffic environments many messages can arrive during a single poll loop, therefore syslog-ng reads only a fixed number of messages from each source. The **Messages fetched in a single poll** option specifies the number of messages read during a poll loop from a single source.

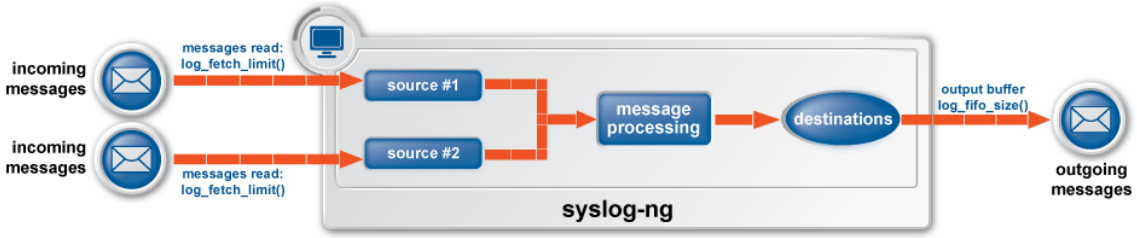

Figure 2.3. Managing log messages in syslog-ng

Note
The **Messages fetched in a single poll** option of SSB can be set as a global option at **Log > Options**.

Every destination has its own output buffer. The output buffer is needed because the destination might not be able to accept all messages immediately. On SSB, the **Output memory buffer** parameter sets the size of the output buffer. The output buffer must be larger than the **Messages fetched in a single poll** of the sources, to ensure that every message read during the poll loop fits into the output buffer. If the log path sends messages to a destination from multiple sources, the output buffer must be large enough to store the incoming messages of every source.

TCP and TLS sources can receive the logs from several incoming connections (for example many different clients or applications). For such sources, syslog-ng reads messages from every connection, thus the **Messages fetched in a single poll** parameter applies individually to every connection of the source.

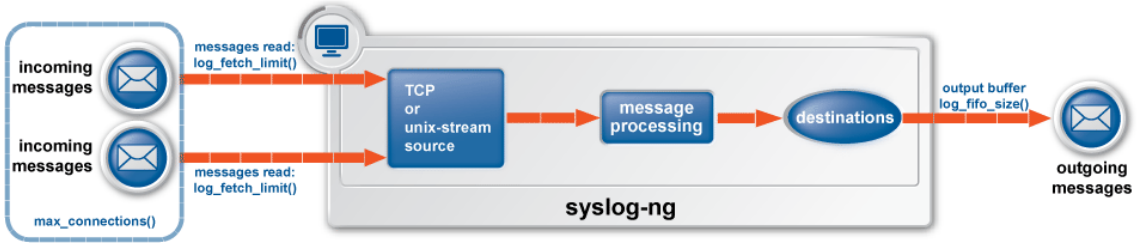

Figure 2.4. Managing log messages of TCP sources in syslog-ng

The flow-control of syslog-ng introduces a control window to the source that tracks how many messages can syslog-ng accept from the source. Every message that syslog-ng reads from the source decreases the number of free slots by one; every message that syslog-ng successfully sends from the output buffer increases the number of free slots by one. If the window is full (that is, there are no free slots), syslog-ng stops reading messages from the source. The initial size of the control window is by default *100*: the **Output memory buffer** must be larger than this value in order for flow-control to have any effect. If a source accepts messages from multiple connections, all messages use the same control window.

When flow-control is used, every source has its own control window. As a worst-case situation, the output buffer of the destination must be set to accommodate all messages of every control window, that is, the **Output memory buffer** of the destination must be greater than **<Number of sources>*<Initial window size>**. This applies to every source that sends logs to the particular destination, thus if two sources having several connections

and heavy traffic send logs to the same destination, the control windows of every source must fit into the output buffer of the destination. Otherwise, syslog-ng does not activate the flow-control, and messages may be lost.

2.3.1. Flow-control and multiple destinations

Using flow-control on a source has an important side-effect if the messages of the source are sent to multiple destinations. If flow-control is in use and one of the destinations cannot accept the messages, the other destinations do not receive any messages either, because syslog-ng stops reading the source. For example, if messages from a source are sent to a remote server and also stored locally in a file, and the network connection to the server becomes unavailable, neither the remote server nor the local file will receive any messages. This side-effect of the flow-control can be avoided by using the disk-based buffering feature of syslog-ng.

Note
Creating separate log paths for the destinations that use the same flow-controlled source does not help avoiding the problem.

2.4. Receiving logs from a secure channel

The syslog-ng Store Box receive log messages securely over the network using the Transport Layer Security (TLS) protocol (TLS is an encryption protocol over the TCP/IP network protocol).

TLS uses certificates to authenticate and encrypt the communication, as illustrated on the following figure:

CLIENT
SERVER
1. THE CLIENT CONNECTS TO THE SERVER
CA2 CERTIFICATE
2. THE SERVER SENDS THE SERVER CERTIFICATE
SERVER CERTIFICATE Issued by CA2
3. THE CLIENT VERIFIES THE SERVER CERTIFICATE USING THE CA2 CERIFICATE
CLIENT CERTIFICATE Issued by CA1
4. THE CLIENT SENDS THE CLIENT CERTIFICATE TO THE SERVER
CA1 CERTIFICATE
5. THE SERVER VERIFIES THE CLIENT CERTIFICATE USING THE CA1 CERTIFICATE

Figure 2.5. Certificate-based authentication

The client sending the logs authenticates SSB by requesting its certificate and public key. Optionally, SSB can also request a certificate from the client, thus mutual authentication is also possible.

In order to use TLS encryption in syslog-ng, the following elements are required:

- A certificate on SSB that identifies SSB. This is available by default.
- The certificate of the Certificate Authority that issued the certificate of SSB must be available on the syslog-ng client.

When using mutual authentication to verify the identity of the clients, the following elements are required:

- A certificate must be available on the syslog-ng client. This certificate identifies the syslog-ng client.
- The certificate of the Certificate Authority that issued the certificate of the syslog-ng client must be available on SSB.

Mutual authentication ensures that SSB accepts log messages only from authorized clients.

For details on configuring TLS communication in syslog-ng, see *Chapter 7, Configuring message sources (p. 126)*.

2.5. Reliable Log Transfer Protocol™

The SSB application can receive log messages in a reliable way over the TCP transport layer using the Reliable Log Transfer Protocol™ (RLTP™). RLTP™ is a proprietary transport protocol that prevents message loss during connection breaks. The transport is used between syslog-ng Premium Edition hosts and SSB (for example, a client and SSB, or a client-relay-SSB), and interoperates with the flow-control and reliable disk-buffer mechanisms of syslog-ng Premium Edition, thus providing the best way to prevent message loss. The sender detects which messages has the receiver successfully received. If messages are lost during the transfer, the sender resends the missing messages, starting from the last successfully received message. Therefore, messages are not duplicated at the receiving end in case of a connection break (however, in failover mode this is not completely ensured). RLTP™ also allows to receive encrypted connections.

2.6. Network interfaces

The SSB hardware has five network interfaces: the external, the management, the internal currently not used in SSB, the HA, and the IPMI interface. For details on hardware installation, see *Appendix B, syslog-ng Store Box Hardware Installation Guide (p. 238)*.

The *external* interface is used for communication between SSB and the clients: clients send the syslog messages to the external interface of SSB. Also, the initial configuration of SSB is always performed using the external interface (For details on the initial configuration, see *Procedure 3.2, Configuring SSB with the Welcome Wizard (p. 22)*). The external interface is used for management purposes if the management interface is not configured. The external interface uses the Ethernet connector labeled as `1` (or `EXT`).

The *management* interface is used exclusively for communication between SSB and the auditors or the administrators of SSB. Incoming connections are accepted only to access the SSB web interface, other connections targeting this interface are rejected. The management interface uses the Ethernet connector labeled as `2` (or `MGMT`).

The routing rules determine which interface is used for transferring remote backups and syslog messages of SSB.

Tip
It is recommended to direct backups, syslog and SNMP messages, and e-mail alerts to the management interface. For details, see *Procedure 4.3.2, Routing management traffic to the management interface (p. 39)*.

If the management interface is not configured, the external interface takes the role of the management interface.

The *high availability* interface (*HA*) is an interface reserved for communication between the nodes of SSB clusters. The HA interface uses the Ethernet connector labeled as *4* (or *HA*). For details on high availability, see *Section 2.7, High Availability support in SSB (p. 9)*.

The Intelligent Platform Management Interface (*IPMI*) interface allows system administrators to monitor system health and to manage SSB events remotely. IPMI operates independently of the operating system of SSB.

2.7. High Availability support in SSB

High availability clusters can stretch across long distances, such as nodes across buildings, cities or even continents. The goal of HA clusters is to support enterprise business continuity by providing location-independent load balancing and failover.

In high availability (HA) mode two SSB units (called master and slave nodes) having identical configuration are operating simultaneously. The master shares all data with the slave node, and if the master node stops functioning, the other one becomes immediately active, so the servers are continuously accessible. The slave node takes over the MAC addresses of the interfaces of the master node.

You can find more information on managing a high availability SSB cluster in *Section 6.2, Managing a high availability SSB cluster (p. 85)*.

2.8. Firmware in SSB

The SSB firmware is separated into two parts: an *external* and an *internal* firmware.

- The *external* firmware (also called boot firmware) boots up SSB, provides the high availability support, and starts the internal firmware. The external firmware changes very rarely.
- The *internal* firmware (also called core firmware) handles everything else: provides the web interface, receives and processes log messages and so on. The internal firmware is updated regularly as new features are added to SSB.

Both firmwares can be updated from the SSB web interface. For details, see *Section 6.3, Upgrading SSB (p. 94)*.

2.8.1. Firmwares and high availability

When powering on the SSB nodes in high availability mode, both nodes boot and start the boot firmware. The boot firmwares then determine which unit is the master: the core firmware is started only on the master node.

Upgrading the SSB firmware via the web interface automatically upgrades the firmware on both nodes.

2.9. Versions and releases of SSB

As of June 2011, the following release policy applies to syslog-ng Store Box:

- *Long Term Supported or LTS releases* (for example, SSB 3 LTS) are supported for 3 years after their original publication date and for 1 year after the next LTS release is published (whichever date is later). The second digit of the revisions of such releases is 0 (for example, SSB 3.0.1). Maintenance releases to LTS releases contain only bugfixes and security updates.

- *Feature releases* (for example, SSB 3 F1) are supported for 6 months after their original publication date and for 2 months after succeeding Feature or LTS Release is published (whichever date is later). Feature releases contain enhancements and new features, presumably 1-3 new feature per release. Only the last feature release is supported (for example when a new feature release comes out, the last one becomes unsupported within two months).

For a full description on stable and feature releases, see *Stable and feature releases*.

Warning
Downgrading from a feature release is not supported. If you upgrade from an LTS release (for example, 3.0) to a feature release (3.1), you have to keep upgrading with each new feature release until the next LTS version (in this case, 4.0) is published.

2.10. Licenses

SSB's license determines the number of individual hosts (also called log source hosts) that can send log messages to SSB.

A log source host is a host or network device (including syslog-ng clients and relays) that sends logs to the syslog-ng server. Log source hosts can be servers, routers, desktop computers, or other devices capable of sending syslog messages or running syslog-ng. Log source hosts are identified by their IP addresses, so virtual machines and vhosts are separately counted.

The SSB license also allows you to download the syslog-ng Premium Edition application (including the syslog-ng Agent for Windows) and install it on any supported platform to use it as a log collector agent for SSB.

Contact BalaBit or your local distributor for details. For details on installing a new license, see *Procedure 6.3.6, Updating the SSB license (p. 99)*.

2.11. The structure of a log message

The following sections describe the structure of log messages. Currently there are two standard syslog message formats:

- The old standard described in RFC 3164 (also called the BSD-syslog or the legacy-syslog protocol): see *Section 2.11.1, BSD-syslog or legacy-syslog messages (p. 10)*
- The new standard described in RFC 5424 (also called the IETF-syslog protocol): see *Section 2.11.2, IETF-syslog messages (p. 12)*

2.11.1. BSD-syslog or legacy-syslog messages

This section describes the format of a syslog message, according to the legacy-syslog or BSD-syslog protocol (see *RFC 3164*). A syslog message consists of the following parts:

- *PRI*
- *HEADER*
- *MSG*

The total message must be shorter than 1024 bytes.

The following is a sample syslog message: `<133>Feb 25 14:09:07 webserver syslogd: restart`. The message corresponds to the following format: `<priority>timestamp hostname application: message`. The different parts of the message are explained in the following sections.

Note
The syslog-ng application supports longer messages as well. For details, see the **Message size** option. However, it is not recommended to enable messages larger than the packet size when using UDP destinations.

2.11.1.1. The PRI message part

The PRI part of the syslog message (known as Priority value) represents the Facility and Severity of the message. Facility represents the part of the system sending the message, while severity marks its importance. The Priority value is calculated by first multiplying the Facility number by 8 and then adding the numerical value of the Severity. The possible facility and severity values are presented below.

Note
Facility codes may slightly vary between different platforms. The syslog-ng application accepts facility codes as numerical values as well.

The following table lists the facility values.

Numerical Code	Facility
0	kernel messages
1	user-level messages
2	mail system
3	system daemons
4	security/authorization messages
5	messages generated internally by syslogd
6	line printer subsystem
7	network news subsystem
8	UUCP subsystem
9	clock daemon
10	security/authorization messages
11	FTP daemon
12	NTP subsystem
13	log audit
14	log alert

Numerical Code	Facility
15	clock daemon
16-23	locally used facilities (local0-local7)

Table 2.1. syslog Message Facilities

The following table lists the severity values.

Numerical Code	Severity
0	Emergency: system is unusable
1	Alert: action must be taken immediately
2	Critical: critical conditions
3	Error: error conditions
4	Warning: warning conditions
5	Notice: normal but significant condition
6	Informational: informational messages
7	Debug: debug-level messages

Table 2.2. syslog Message Severities

2.11.1.2. The HEADER message part

The HEADER part contains a timestamp and the hostname (without the domain name) or the IP address of the device. The timestamp field is the local time in the `Mmm dd hh:mm:ss` format, where:

- *Mmm* is the English abbreviation of the month: Jan, Feb, Mar, Apr, May, Jun, Jul, Aug, Sep, Oct, Nov, Dec.
- *dd* is the day of the month on two digits. If the day of the month is less than 10, the first digit is replaced with a space. (For example `Aug 7`.)
- *hh:mm:ss* is the local time. The hour (hh) is represented in a 24-hour format. Valid entries are between 00 and 23, inclusive. The minute (mm) and second (ss) entries are between 00 and 59 inclusive.

2.11.1.3. The MSG message part

The MSG part contains the name of the program or process that generated the message, and the text of the message itself. The MSG part is usually in the following format: `program[pid]: message text`.

2.11.2. IETF-syslog messages

This section describes the format of a syslog message, according to the IETF-syslog protocol (see *RFC 5424-5428*).A syslog message consists of the following parts:

- `HEADER (includes the PRI as well)`
- `STRUCTURED-DATA`
- `MSG`

The following is a sample syslog message:

```
<34>1 2003-10-11T22:14:15.003Z mymachine.example.com su - ID47 - BOM'su root' failed
 for lonvick on /dev/pts/8
```

The message corresponds to the following format:

```
<priority>VERSION ISOTIMESTAMP HOSTNAME APPLICATION PID MESSAGEID STRUCTURED-DATA
 MSG
```

In this example, the Facility has the value of 4, severity is 2, so PRI is 34. The VERSION is 1. The message was created on 11 October 2003 at 10:14:15pm UTC, 3 milliseconds into the next second. The message originated from a host that identifies itself as "mymachine.example.com". The APP-NAME is "su" and the PROCID is unknown. The MSGID is "ID47". The MSG is "'su root' failed for lonvick...", encoded in UTF-8. The encoding is defined by the BOM. There is no STRUCTURED-DATA present in the message, this is indicated by "-" in the STRUCTURED-DATA field. The MSG is "'su root' failed for lonvick...".

The HEADER part of the message must be in plain ASCII format, the parameter values of the STRUCTURED-DATA part must be in UTF-8, while the MSG part should be in UTF-8. The different parts of the message are explained in the following sections.

2.11.2.1. The PRI message part

The PRI part of the syslog message (known as Priority value) represents the Facility and Severity of the message. Facility represents the part of the system sending the message, while severity marks its importance. The Priority value is calculated by first multiplying the Facility number by 8 and then adding the numerical value of the Severity. The possible facility and severity values are presented below.

Note
Facility codes may slightly vary between different platforms. The syslog-ng application accepts facility codes as numerical values as well.

Numerical Code	Facility
0	kernel messages
1	user-level messages
2	mail system
3	system daemons
4	security/authorization messages
5	messages generated internally by syslogd
6	line printer subsystem
7	network news subsystem
8	UUCP subsystem

Source: http://tools.ietf.org/html/rfc5424

Numerical Code	Facility
9	clock daemon
10	security/authorization messages
11	FTP daemon
12	NTP subsystem
13	log audit
14	log alert
15	clock daemon
16-23	locally used facilities (local0-local7)

Table 2.3. syslog Message Facilities

The following table lists the severity values.

Numerical Code	Severity
0	Emergency: system is unusable
1	Alert: action must be taken immediately
2	Critical: critical conditions
3	Error: error conditions
4	Warning: warning conditions
5	Notice: normal but significant condition
6	Informational: informational messages
7	Debug: debug-level messages

Table 2.4. syslog Message Severities

2.11.2.2. The HEADER message part

The HEADER part contains the following elements:

- *VERSION*: Version number of the syslog protocol standard. Currently this can only be `1`.
- *ISOTIMESTAMP*: The time when the message was generated in the ISO 8601 compatible standard timestamp format (yyyy-mm-ddThh:mm:ss+-ZONE), for example: `2006-06-13T15:58:00.123+01:00`.
- *HOSTNAME*: The machine that originally sent the message.
- *APPLICATION*: The device or application that generated the message
- *PID*: The process name or process ID of the syslog application that sent the message. It is not necessarily the process ID of the application that generated the message.
- *MESSAGEID*: The ID number of the message.

Note
The syslog-ng application supports other timestamp formats as well, like ISO, or the PIX extended format. The timestamp used in the IETF-syslog protocol is derived from RFC3339, which is based on ISO8601. For details, see the `ts_format()` option in *The syslog-ng Premium Edition 5 LTS Administrator Guide*.

2.11.2.3. The STRUCTURED-DATA message part

The STRUCTURED-DATA message part may contain meta- information about the syslog message, or application-specific information such as traffic counters or IP addresses. STRUCTURED-DATA consists of data blocks enclosed in brackets (`[]`). Every block include the ID of the block, and one or more `name=value` pairs. The syslog-ng application automatically parses the STRUCTURED-DATA part of syslog messages, which can be referenced in macros (see *The syslog-ng Premium Edition 5 LTS Administrator Guide* for details). An example STRUCTURED-DATA block looks like:

```
[exampleSDID@0 iut="3" eventSource="Application" eventID="1011"][examplePriority@0
 class="high"]
```

2.11.2.4. The MSG message part

The MSG part contains the text of the message itself. The encoding of the text must be UTF-8 if the BOM character is present in the message. If the message does not contain the BOM character, the encoding is treated as unknown. Usually messages arriving from legacy sources do not include the BOM character.

Chapter 3. The Welcome Wizard and the first login

This chapter describes the initial steps of configuring SSB. Before completing the steps below, unpack, assemble, and power on the hardware. Connect at least the external network interface to the local network, or directly to the computer from which SSB will be configured.

Note
For details on unpacking and assembling the hardware, see *Appendix B, syslog-ng Store Box Hardware Installation Guide (p. 238)*. For details on how to create a high availability SSB cluster, see *Procedure B.2, Installing two SSB units in HA mode (p. 240)*.

3.1. The initial connection to SSB

SSB can be connected from a client machine using any modern web browser.

Note
For details on supported browsers, see *Section 4.1, Supported web browsers and operating systems (p. 31)*

SSB can be accessed from the local network. Starting with version 2.1, SSB attempts to receive an IP address automatically via DHCP. If it fails to obtain an automatic IP address, it starts listening for HTTPS connections on the `192.168.1.1` IP address. Note that certain switch configurations and security settings can interfere with SSB receiving an IP address via DHCP. SSB accepts connections via its *external* interface (`EXT`, for details on the network interfaces, see *Section 2.6, Network interfaces (p. 8)*).

Tip
The SSB console displays the IP address the external interface is listening on.

If SSB is listening on the `192.168.1.1` address, note that the `192.168.1.0/24` subnet must be accessible from the client. If the client machine is in a different subnet (for example its IP address is `192.168.10.X`), but in the same network segment, the easiest way is to assign an alias IP address to the client machine. Creating an alias IP on the client machine virtually puts both the client and SSB into the same subnet, so that they can communicate. To create an alias IP complete the following steps.

- For details on creating an alias IP on Microsoft Windows, see *Procedure 3.1.1, Creating an alias IP address (Microsoft Windows) (p. 17)*.
- For details on creating an alias IP on Linux, see *Procedure 3.1.2, Creating an alias IP address (Linux) (p. 20)*.

- If configuring an alias interface is not an option for some reason, you can modify the IP address of SSB. For details, see *Procedure 3.1.3, Modifying the IP address of SSB (p. 21)*.

Warning
The Welcome Wizard can be accessed only using the external network interface of SSB, as the management interface is not configured yet.

3.1.1. Procedure – Creating an alias IP address (Microsoft Windows)

Purpose:

This procedure describes how to assign an alias IP address to a network interface on Microsoft Windows platforms.

Steps:

Step 1. Navigate to **Start menu > Settings > Network Connections**.

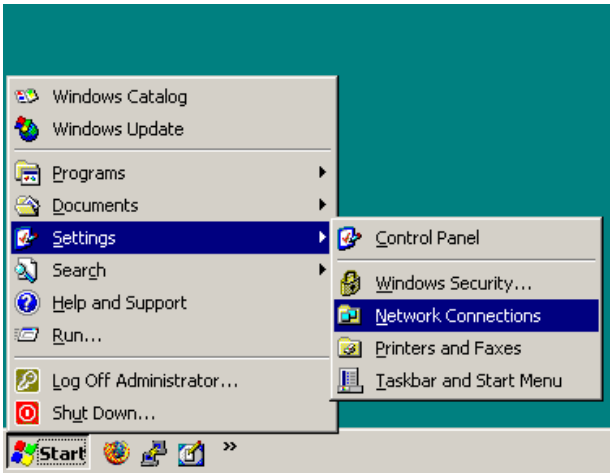

Figure 3.1.

Step 2. Double click on the **Local Area Connection** and then click **Properties**.

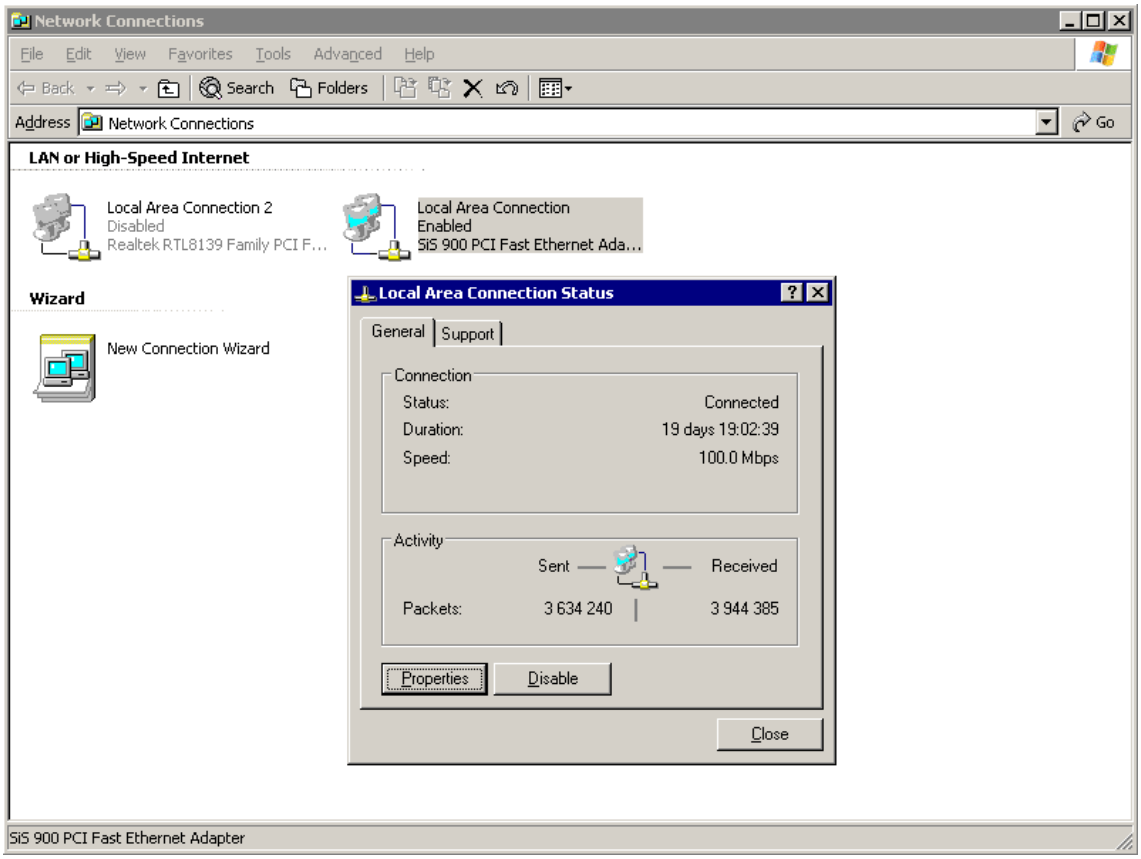

Figure 3.2.

Step 3. Select the **Internet Protocol (TCP/IP)** component in the list and click **Properties**.

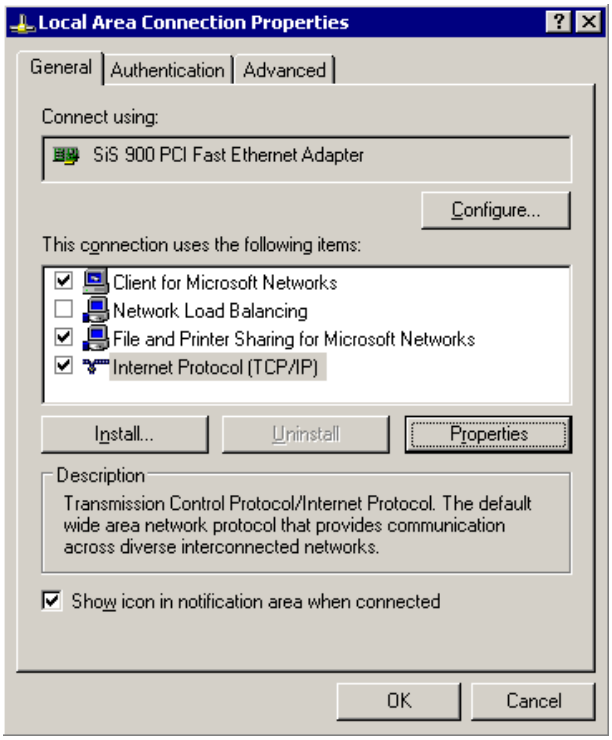

Figure 3.3.

Step 4. To display the Advanced TCP/IP Settings window, click **Advanced**.

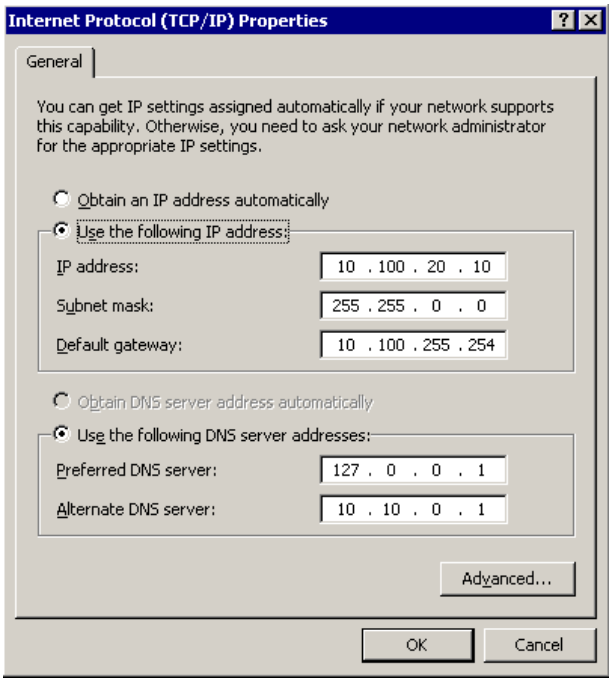
Internet Protocol (TCP/IP) Properties
General
You can get IP settings assigned automatically if your network supports this capability. Otherwise, you need to ask your network administrator for the appropriate IP settings.
Obtain an IP address automatically
Use the following IP address:
IP address:
10 . 100 . 20 . 10
Subnet mask:
255 . 255 . 0 . 0
Default gateway:
10 . 100 . 255 . 254
Obtain DNS server address automatically
Use the following DNS server addresses:
Preferred DNS server:
127 . 0 . 0 . 1
Alternate DNS server:
10 . 10 . 0 . 1
Advanced...
OK
Cancel

Figure 3.4.

Step 5. Select the **IP Settings** tab and in the **IP Addresses** section, click **Add**.

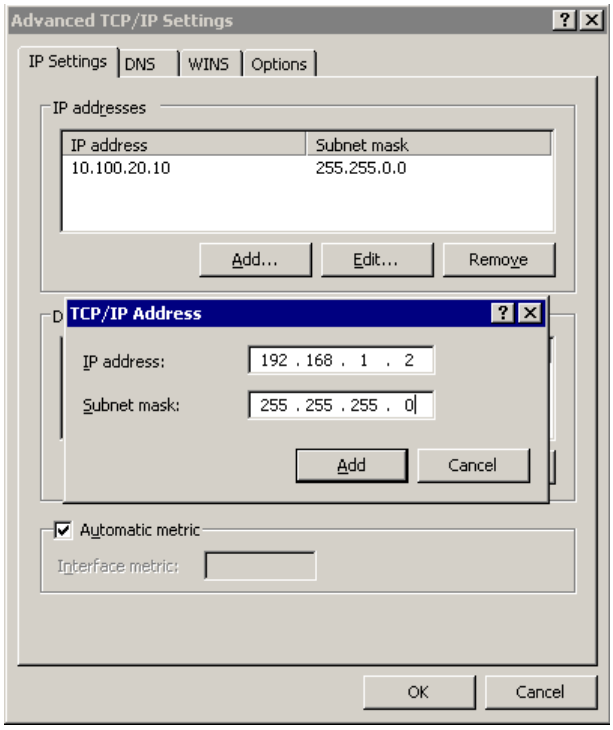

Figure 3.5.

Step 6. Into the **IP Address** field, enter *192.168.1.2*. Into the **Netmask** field, enter *255.255.255.0*.

Warning
If your internal network uses the *192.168.1.0/24* IP range, the *192.168.1.1* and *192.168.1.2* addresses might already be in use. In this case, disconnect SSB from the network, and connect directly a computer to its external interface using a standard cross-link cable.

Step 7. To complete the procedure, click **Add** .

3.1.2. Procedure – Creating an alias IP address (Linux)

Purpose:

This procedure describes how to assign an alias IP address to a network interface on Linux platforms.

Steps:

Step 1. Start a terminal console (for example `gnome-terminal`, `konsole`, `xterm`, and so on).

Step 2. Issue the following command as root:

```
ifconfig <ethX>:0 192.168.1.2
```

where <*ethX*> is the ID of the network interface of the client, usually *eth0* or *eth1*.

Step 3. Issue the `ifconfig` command. The *`<ethX>:0`* interface appears in the output, having *`inet addr:192.168.1.2`* .

Step 4. Issue the `ping -c 3 192.168.1.1` command to verify that SSB is accessible. A similar result is displayed:

```
user@computer:~$ ping -c 3 192.168.1.1
                   PING 192.168.1.1 (192.168.1.1) 56(84) bytes of data.
                   64 bytes from 192.168.1.1: icmp-seq=1 ttl=63 time=0.357
ms
                   64 bytes from 192.168.1.1: icmp-seq=2 ttl=63 time=0.306
ms
                   64 bytes from 192.168.1.1: icmp-seq=3 ttl=63 time=0.314
ms

                   --- 192.168.1.1 ping statistics ---
                   3 packets transmitted, 3 received, 0% packet loss, time
2013ms
                   rtt min/avg/max/mdev = 0.306/0.325/0.357/0.030 ms
```

Open the page *https://192.168.1.1* from your browser and accept the certificate shown. The Welcome Wizard of SSB appears.

3.1.3. Procedure – Modifying the IP address of SSB

Purpose:

To configure SSB to listen for connections on a custom IP address, complete the following steps.

Warning
Use this procedure only before the initial configuration of SSB, that is, before completing the Welcome Wizard. For details on changing the IP address or other network settings of a configured SSB system, see *Section 4.3, Network settings (p. 36)*.

Steps:

Step 1. Access SSB from the local console, and log in with username *`root`* and password *`default`*.

Step 2. In the Console Menu, select **Shells** > **Core shell**.

Step 3. Change the IP address of SSB:
`ifconfig eth0 <IP-address> netmask 255.255.255.0`

Replace `<IP-address>` with an IPv4 address suitable for your environment.

Step 4. Set the default gateway using the following command:
`route add default gw <IP-of-default-gateway>`

Replace `<IP-of-default-gateway>` with the IP address of the default gateway.

Step 5. Type `exit`, then select **Logout** from the Console Menu.

Step 6. Open the page *https://<IP-address-you-set-for-SSB>* from your browser and accept the certificate shown. The Welcome Wizard of SSB appears.

3.2. Procedure – Configuring SSB with the Welcome Wizard

Purpose:

The Welcome Wizard guides you through the basic configuration steps of SSB. All parameters can be modified before the last step by using the **Back** button of the wizard, or later via the web interface of SSB.

Steps:

Step 1. Open the *https://<IP-address-of-SSB-external-interface>* page in your browser and accept the displayed certificate. The Welcome Wizard of SSB appears.

Tip
The SSB console displays the IP address the external interface is listening on. SSB either receives an IP address automatically via DHCP, or if a DHCP server is not available, listens on the `192.168.1.1` IP address.

Step 2. When configuring SSB for the first time, click **Next**.

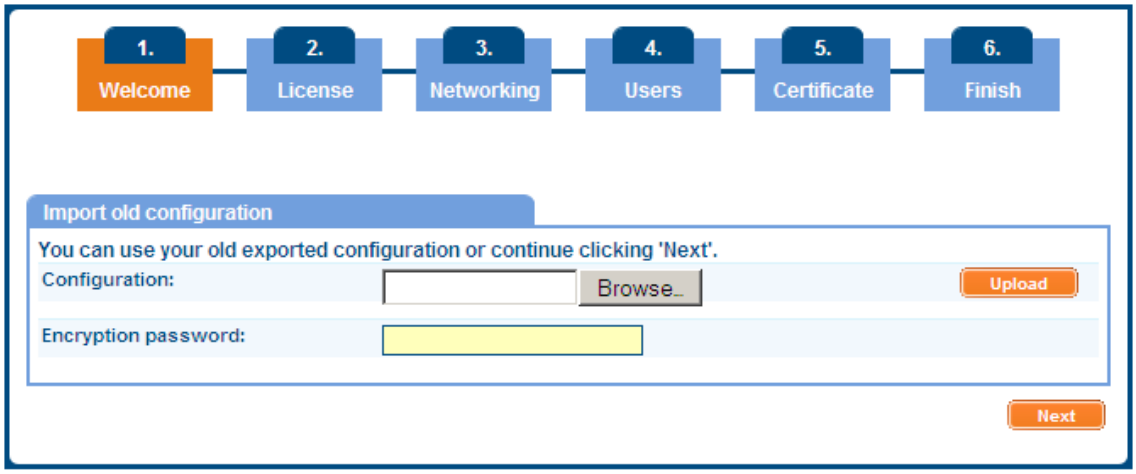

Figure 3.6. The Welcome Wizard

It is also possible to import an existing configuration from a backup file. Use this feature to restore a backup configuration after a recovery, or to migrate an existing SSB configuration to a new device.

Step a. Click **Browse** and select the configuration file to import.

Note
It is not possible to directly import a GPG-encrypted configuration into SSB, it has to be decrypted locally first.

Step b. Enter the passphrase used when the configuration was exported into the **Encryption passphrase** field.

For details on restoring configuration from a configuration backup, see *Procedure 16.7, Restoring SSB configuration and data (p. 235)*

Step c. Click **Import**.

Warning
If you use the Import function to copy a configuration from one SSB to another, do not forget to configure the IP addresses of the second SSB. Having two devices with identical IP addresses on the same network leads to errors.

Step 3. Accept the End User License Agreement and install the SSB license

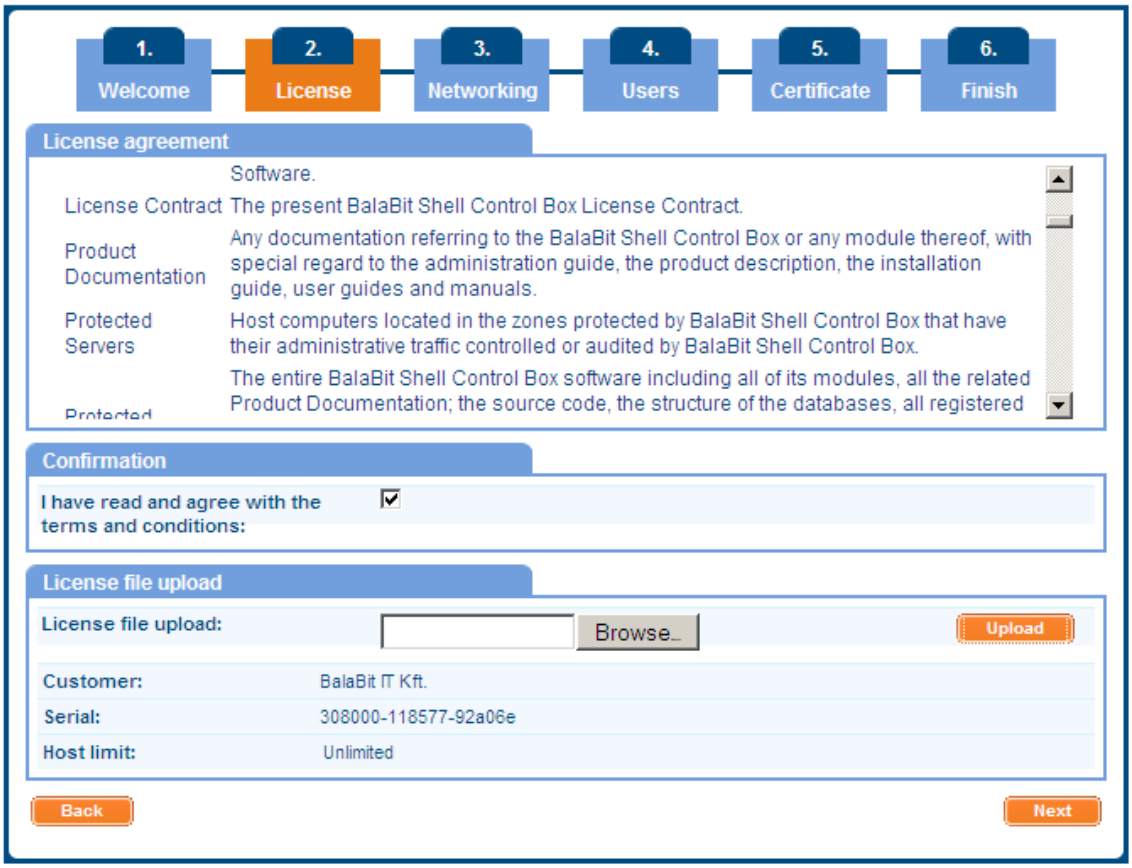
1. Welcome
2. License
3. Networking
4. Users
5. Certificate
6. Finish
License agreement
Software.
License Contract
The present BalaBit Shell Control Box License Contract.
Product Documentation
Any documentation referring to the BalaBit Shell Control Box or any module thereof, with special regard to the administration guide, the product description, the installation guide, user guides and manuals.
Protected Servers
Host computers located in the zones protected by BalaBit Shell Control Box that have their administrative traffic controlled or audited by BalaBit Shell Control Box.
The entire BalaBit Shell Control Box software including all of its modules, all the related Product Documentation; the source code, the structure of the databases, all registered
Confirmation
I have read and agree with the terms and conditions:
License file upload
License file upload:
Browse...
Upload
Customer:
BalaBit IT Kft.
Serial:
308000-118577-92a06e
Host limit:
Unlimited
Back
Next

Figure 3.7. The EULA and the license key

Step a. Read the End User License Agreement and select **Accept**.

Step b. Click **Browse**, select the SSB license file received with SSB, then click **Upload**.

Note
It is not required to manually decompress the license file. Compressed licenses (for example `.zip` archives) can also be uploaded.

Step c. Click **Next**.

Step 4. Fill the fields to configure networking. The meaning of each field is described below. The background of unfilled required fields is red. All parameters can later be modified using the regular interface of SSB.

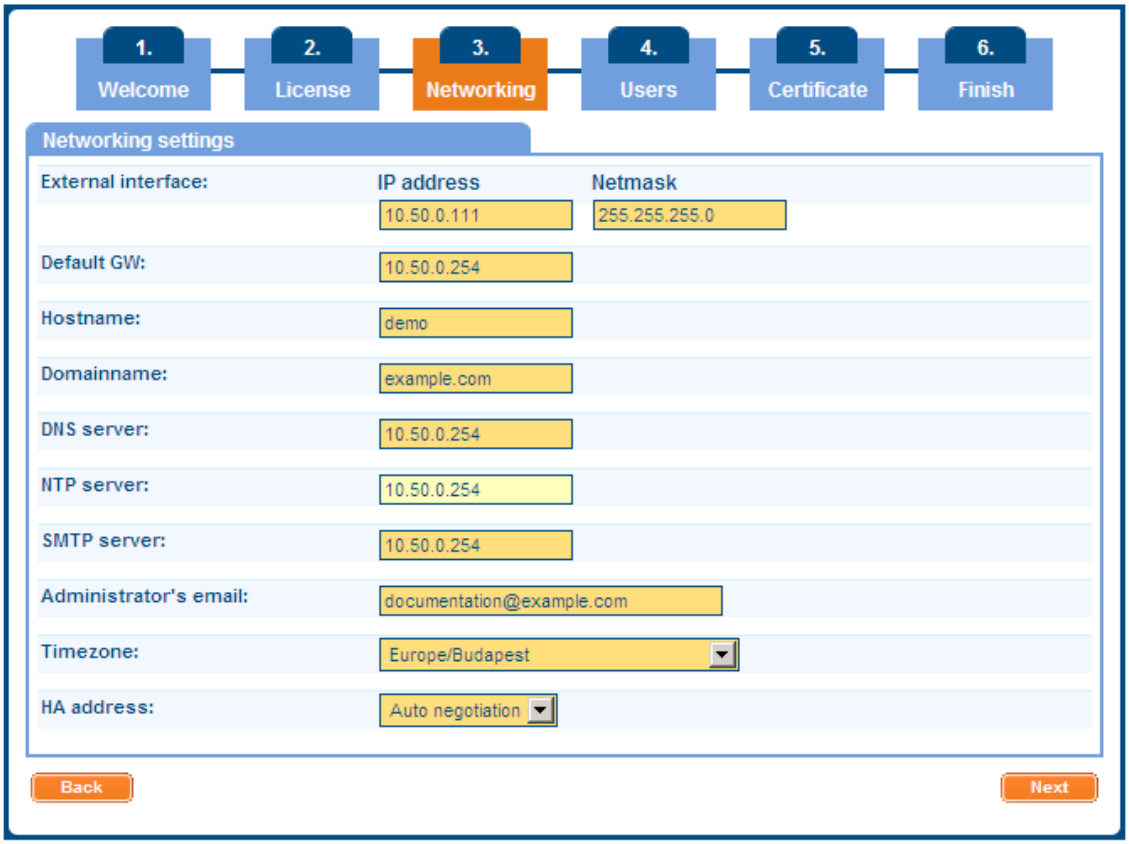

Figure 3.8. Initial networking configuration

Step a. **Hostname**: Name of the machine running SSB (for example *SSB*).

Step b. **Domain name**: Name of the domain used on the network.

Step c. **DNS server**: IP address of the name server used for domain name resolution.

Step d. **NTP server**: The IP address or the hostname of the NTP server.

Step e. **SMTP server**: The IP address or the hostname of the SMTP server used to deliver e-mails.

Step f. **Administrator's e-mail**: E-mail address of the SSB administrator.

Step g. **Timezone**: The timezone where the SSB is located.

Step h. **External interface — IP address**: IP address of the external interface of SSB (for example *192.168.1.1*). The IP address can be chosen from the range of the corresponding physical subnet. Clients will connect the external interface, therefore it must be accessible to them.

Note
Do not use IP addresses that fall into the following ranges:

- `1.2.0.0/16` (reserved for communication between SSB cluster nodes)
- `127.0.0.0/8` (localhost IP addresses)

Step i. **External interface — Netmask**: The IP netmask of the given range in IP format. For example, general class C networks have the 255.255.255.0 netmask.

Step j. **Default gateway**: IP address of the default gateway. When using several network cards, the default gateway is usually in the direction of the external interface.

Step k. **HA address**: The IP address of the high availability (HA) interface. Leave this field on `auto` unless specifically requested by the support team.

Step l. Click **Next**.

Step 5. Enter the passwords used to access SSB.

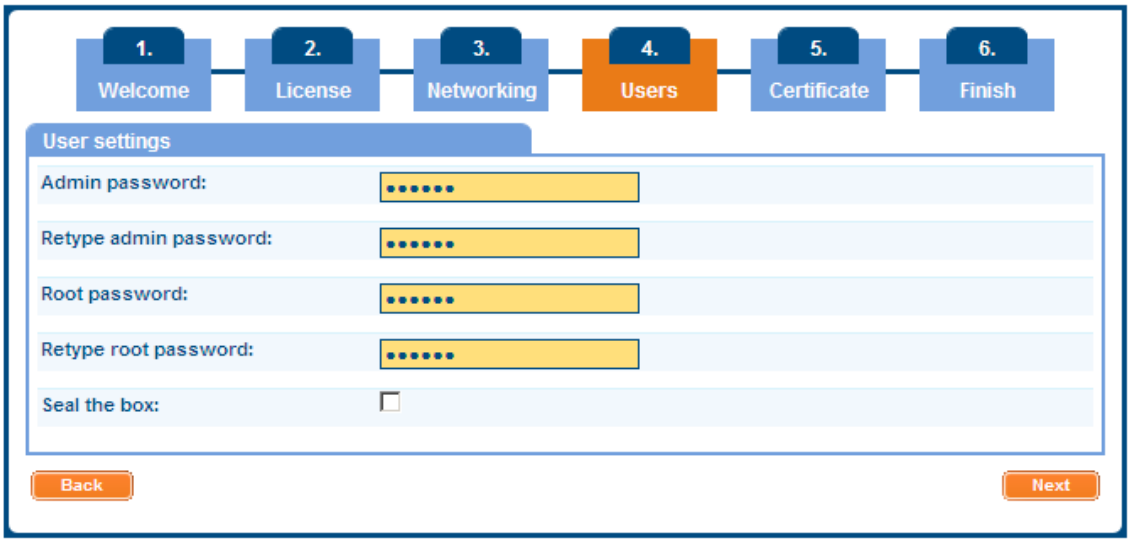

Figure 3.9. Passwords

Note
SSB accepts passwords that are not longer than 150 characters. The following special characters can be used:
`!"#$%&'()*+,-./:;<=>?@[\]^-`{|}`

Step a. **Admin password**: The password of the `admin` user who can access the web interface of SSB.

Step b. **Root password**: The password of the `root` user, required to access SSB via SSH or from the local console.

Note
Accessing SSB using SSH is rarely needed, and recommended only for advanced users for troubleshooting situations.

Step c. If you want to prevent users from accessing SSB remotely via SSH or changing the root password of SSB, select the **Seal the box** checkbox. Sealed mode can be activated later from the web interface as well. For details, see *Section 6.5, Sealed mode (p. 107)*.

Step d. Click **Next**.

Step 6. Upload or create a certificate for the SSB web interface. This SSL certificate will be displayed by SSB to authenticate administrative HTTPS connections to the web interface.

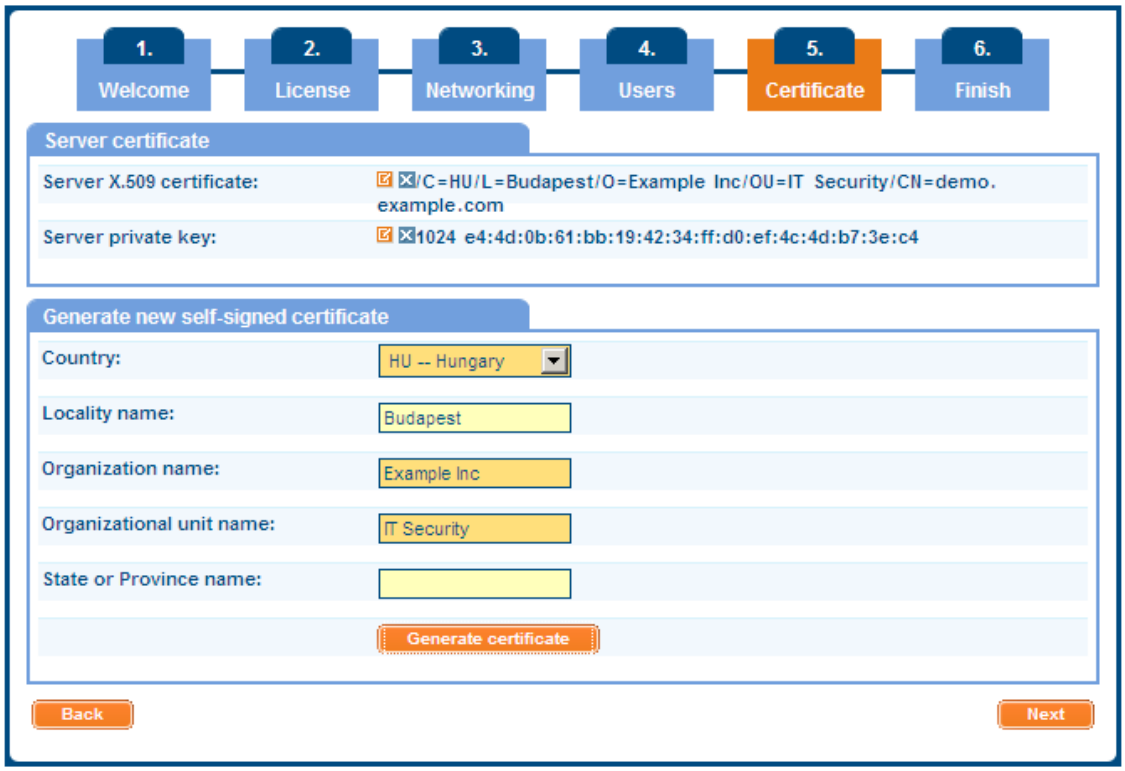

Figure 3.10. Creating a certificate for SSB

To create a self-signed certificate, fill the fields of the **Generate new self-signed certificate** section and click **Generate**. The certificate will be self-signed by the SSB appliance; the hostname of SSB will be used as the issuer and common name.

Step a. **Country**: Select the country where SSB is located (for example HU-Hungary).

Step b. **Locality**: The city where SSB is located (for example Budapest).

Step c. **Organization**: The company who owns SSB (for example Example Inc.).

Step d. **Organization unit**: The division of the company who owns SSB (for example IT Security Department).

Step e. **State or Province**: The state or province where SSB is located.

Step f. Click **Generate**.

If you want to use a certificate that is signed by an external Certificate Authority, in the **Server X.509 certificate** field, click to upload the certificate.

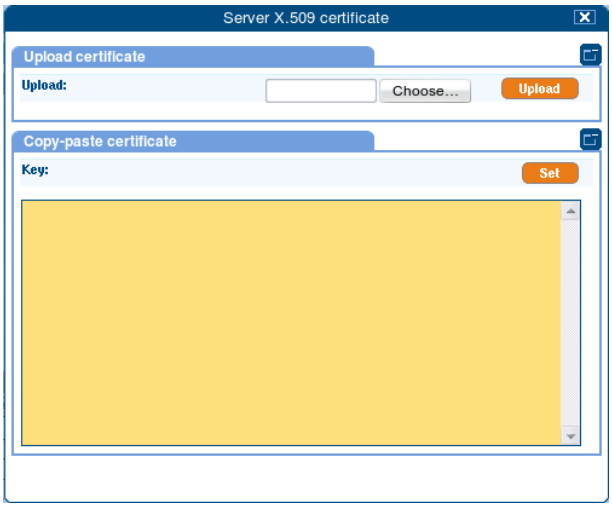

Figure 3.11. Uploading a certificate for SSB

Then in the **Server private key** field click , upload the private key, and enter the password protecting the private key.

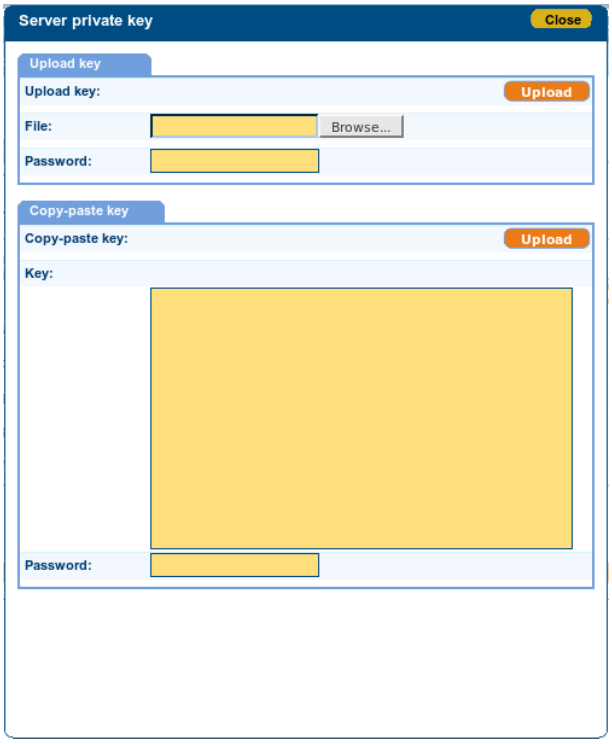

Figure 3.12. Uploading a private key

Note
SSB accepts private keys in PEM (RSA and DSA), PUTTY, and SSHCOM/Tectia format. Password-protected private keys are also supported.

BalaBit recommends using 2048-bit RSA keys (or stronger).

Note
SSB accepts passwords that are not longer than 150 characters. The following special characters can be used: `!"#$%&'()*+,-./:;<=>?@[\]^-`{|}`

Step 7. Review the data entered in the previous steps. This page also displays the certificate generated in the last step; the RSA SSH key of SSB, and information about the license file.

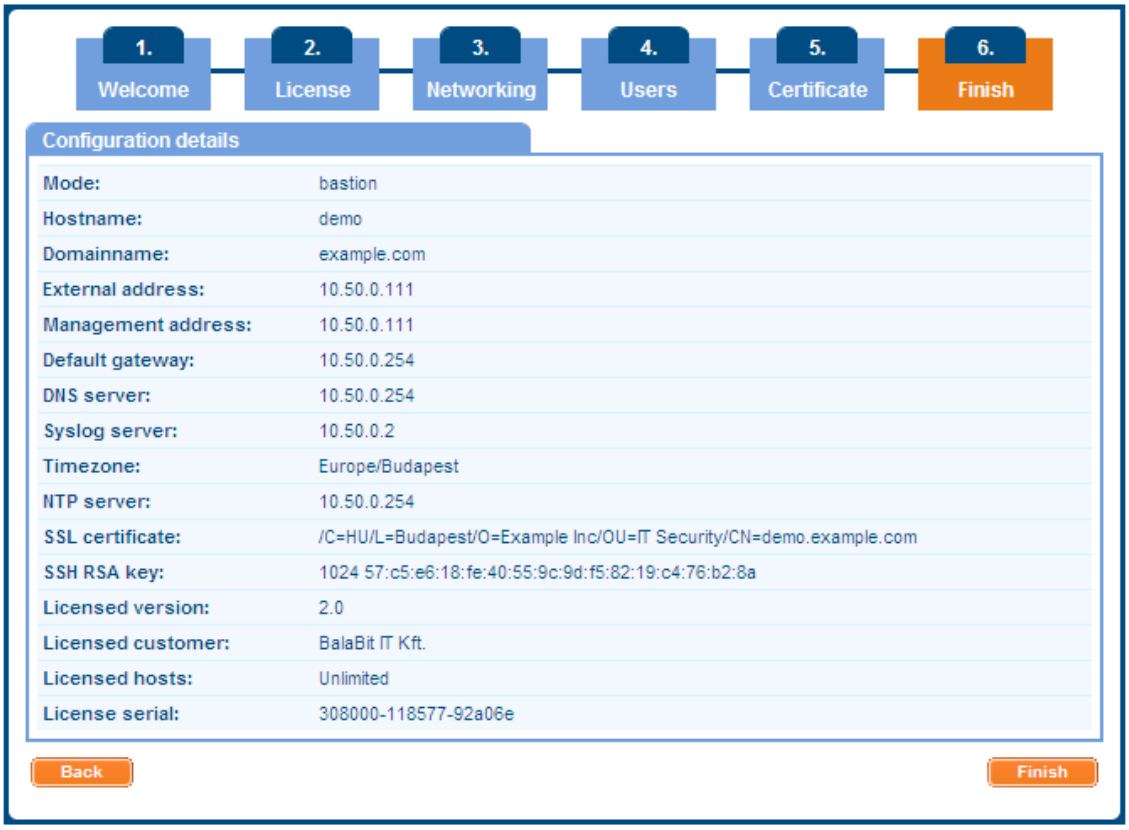

Figure 3.13. Review configuration data

If all information is correct, click **Finish**.

Warning
The configuration takes effect immediately after clicking **Finish**. Incorrect network configuration data can render SSB unaccessible.

SSB is now accessible from the regular web interface via the IP address of its external interface.

Step 8. Your browser is automatically redirected to the IP address set as the external interface of SSB, where you can login to the web interface of SSB using the *admin* username and the password you set for this user in the Welcome Wizard.

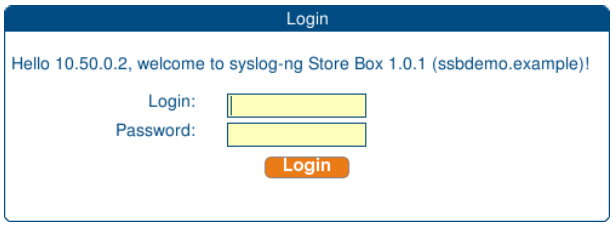

Figure 3.14. Logging in to SSB

Chapter 4. Basic settings

SSB is configured via the web interface. Configuration changes take effect automatically after clicking Commit. Only the modifications of the current page or tab are activated — each page and tab must be committed separately.

- For the list of supported browsers, see *Section 4.1, Supported web browsers and operating systems (p. 31)*.
- For a description of the web interface of SSB, see *Section 4.2, The structure of the web interface (p. 31)*.
- To configure network settings, see *Section 4.3, Network settings (p. 36)*.
- To configure date and time settings, see *Section 4.4, Date and time configuration (p. 40)*.
- To configure system logging and e-mail alerts, see *Section 4.5, SNMP and e-mail alerts (p. 41)*
- To configure system monitoring, see *Section 4.6, Configuring system monitoring on SSB (p. 47)*.
- To configure data and configuration backups, see *Section 4.7, Data and configuration backups (p. 55)*.
- To configure archiving and clean-up, see *Section 4.8, Archiving and cleanup (p. 59)*.
- For a description of the backup and archiving protocols, see *Section 4.9, Parameters of the backup and archiving protocols (p. 62)*.

4.1. Supported web browsers and operating systems

The SSB web interface can be accessed only using TLS encryption and strong cipher algorithms.

Supported browsers: Mozilla Firefox 17 and Microsoft Internet Explorer 8 and 9. The browser must support HTTPS connections, JavaScript, and cookies. Make sure that both JavaScript and cookies are enabled.

When using Internet Explorer 8.0 on Windows Server 2008 R2 Enterprise, disable the Content Advisor in the Internet Explorer. To accomplish this, select **Tools** > **Internet Options** > **Content** > **Content Advisor** > **Disable**.

Other tested browsers: Mozilla Firefox 3.6 and Google Chrome 17.

Note
SSB displays a warning message if your browser is not supported or JavaScript is disabled.

Supported operating systems: Microsoft Windows XP, Windows 2003 Server, Windows Vista, Windows 2008 Server, Windows 7, and Linux.

4.2. The structure of the web interface

The web interface consists of the following main sections:

Main menu: Each menu item displays its options in the main workspace on one or more tabs. Click in front of a main menu item to display the list of available tabs.

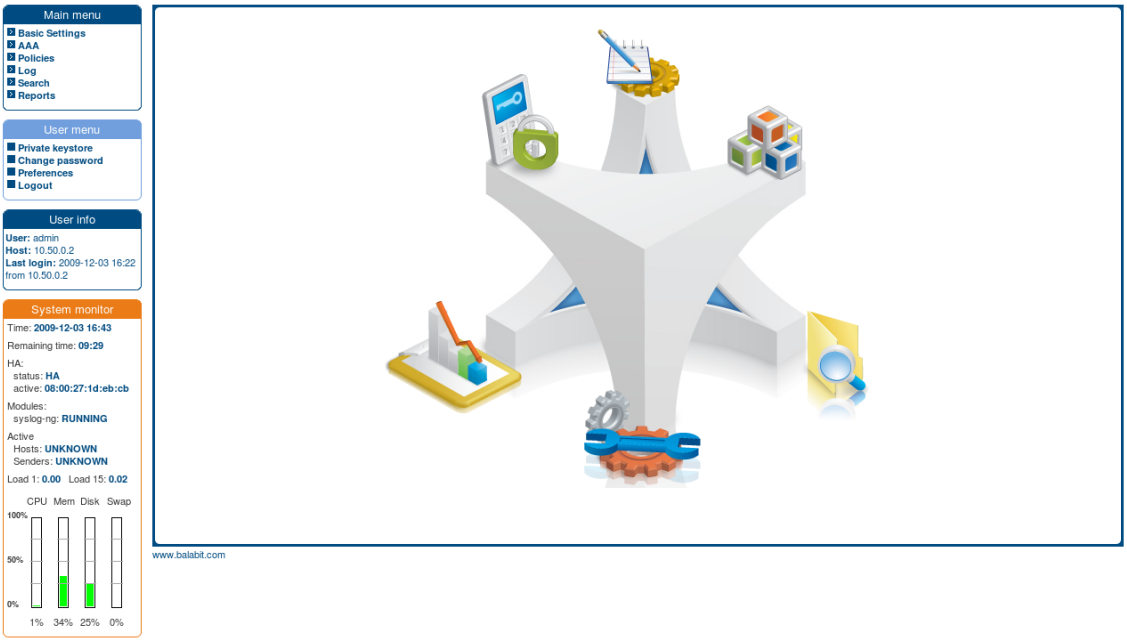

Figure 4.1. Structure of the web interface

User menu: Provides possibilities to change your SSB password; to log out; and disable confirmation dialogs and tooltips using the **Preferences** option.

User info: Provides information about the user currently logged in:

- **User:** username
- **Host:** IP address of the user's computer
- **Last login:** date and IP address of the user's last login

Figure 4.2. User menu and user info

System monitor: Displays accessibility and system health information about SSB, including the following:

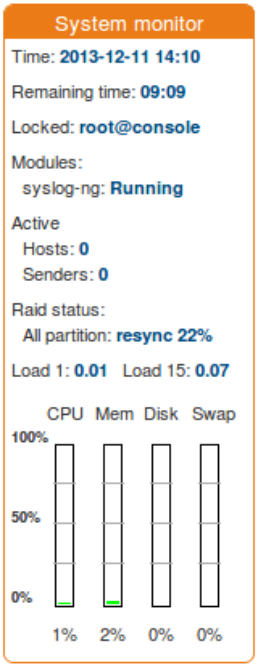

Figure 4.3. System monitor

- **Time:** System date and time.
- **Remaining time:** The time remaining before the session to the web interface times out.

Note
To change timeout settings, navigate to **Basic Settings** > **Management** > **Web interface timeout** and enter the timeout value in minutes.

- **Locked:** Indicates that the interface is locked by another administrator (for details, see *Section 4.2.2, Multiple web users and locking (p. 35)*)
- **Modules:** The status of syslog-ng running on SSB (ideally it is `RUNNING`).
- **License:** License information if the license is not valid, or an evaluation version license has expired.
- **Raid status:** The status of the RAID devices, if synchronization between the disks is in progress.
- **Active:**
  - **Hosts:** the number of clients (log source hosts) where the log messages originate from (for example computers)
  - **Senders:** the number of senders where the log messages directly come from (for example relays)

Example 4.1. Number of hosts and senders
For example: if 300 clients all send log messages directly to SSB the Hosts and Senders are both 300.

If the 300 clients send the messages to 3 relays (assuming that the relays do not send messages themselves) and only the relays communicate directly with SSB then Hosts is 300, while Senders is 3 (the 3 relays).

If the relays also send messages, then Hosts is 303, while Senders is 3 (the 3 relays).

- **HA:**The HA status and the ID of the active node if two SSB units are running in a High Availability cluster. If there are redundant Heartbeat interfaces configured, their status is displayed as well. If the nodes of the cluster are synchronizing data between each other, the progress and the time remaining from the synchronization process is also displayed.

- Average system load during the
  - **Load 1:** last minute
  - **Load 15:** last fifteen minutes

- CPU, memory, hard disk, and swap use. Hover the mouse above the graphical bars to receive a more details in a tooltip, or navigate to **Basic Settings** > **Dashboard** for detailed reports.

The System monitor displays current information about the state of SSB. To display a history of these parameters, go to **Basic Settings** > **Dashboard**. For details, see *Section 16.5, Status history and statistics (p. 226)*.

4.2.1. Elements of the main workspace

The main workspace displays the configuration settings related to the selected main menu item grouped into one or more tabs. Related parameters of a tab are organized into labeled groups or sections, marked with blue outline Interfaces.

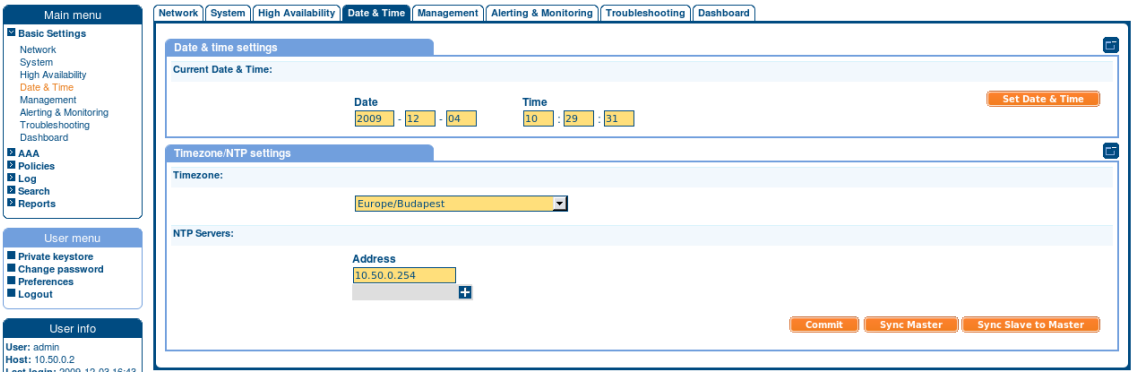

Figure 4.4. Main workspace

- Commit Each page includes one or more orange action buttons. The most common action button is the Commit, which saves and activates the changes of the page.
- /*Show/Hide Details*: Displays or hides additional configuration settings and options.
- , *Create entry*: Create a new row or entry (for example an IP address or a policy).

- , *Delete entry*: Delete a row or an entry (for example an IP address or a policy).
- , *Open/collapse lists*: Open or close a list of options (for example the list of available reports).
- *Modify entries or upload files*: Edit an entry (for example a host key, a list, and so on), or upload a file (for example a private key). These actions open a popup window where the actual modification can be performed.
- , *Position an item in a list*: Modify the order of items in a list. The order of items in a list (for example the order of connections, permitted channels in a channel policy, and so on) is important because when SSB is looking for a policy, it evaluates the list from top to down, and selects the first item completely matching the search criteria. For example, when a client initiates a connection to a protected server, SSB selects the first connection policy matching the client's IP address, the server's IP address, and the target port (the From, To, and Port fields of the connection).

Message window: This popup window displays the responses of SSB to the user's actions, for example **Configuration saved successfully**. Error messages are also displayed here. All messages are included in the system log. For detailed system logs (including message history), see the ***Troubleshooting*** tab of the Basic menu. To make the window appear only for failed actions, navigate to **User menu** > **Preferences** and enable the **Autoclose successful commit messages** option.

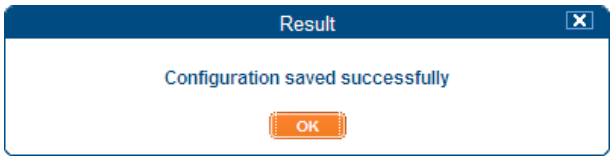

Figure 4.5. Message window

4.2.2. Multiple web users and locking

Multiple administrators can access the SSB web interface simultaneously, but only one of them can modify the configuration. This means that the configuration of SSB is automatically locked when the first administrator who can modify the configuration accesses a configuration page (for example the **Basic Settings**, **AAA**, or **Logs** menu). The username and IP address of the administrator locking the configuration is displayed in the **System Monitor** field. Other administrators must wait until the locking administrator logs out, or the session of the administrator times out. However, it is possible to access the **Search** and **Reporting** menus, or browse the configuration with only View rights (for details, see *Section 5.6, Managing user rights and usergroups (p. 77)*).

Note
If an administrator logs in to SSB using the local console or a remote SSH connection, access via the web interface is completely blocked. Inactive local and SSH connections timeout just like web connections. For details, see *Section 6.4, Accessing the SSB console (p. 103)*.

4.2.3. Web interface timeout

By default, SSB terminates the web session of a user after ten minutes of inactivity. To change value of this timeout, adjust the **Basic Settings** > **Management** > **Web interface timeout** option.

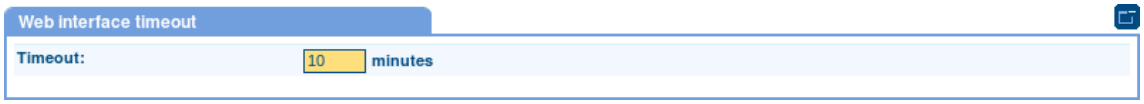

Figure 4.6. Web interface timeout

4.3. Network settings

The **Network** tab contains the network interface and naming settings of SSB.

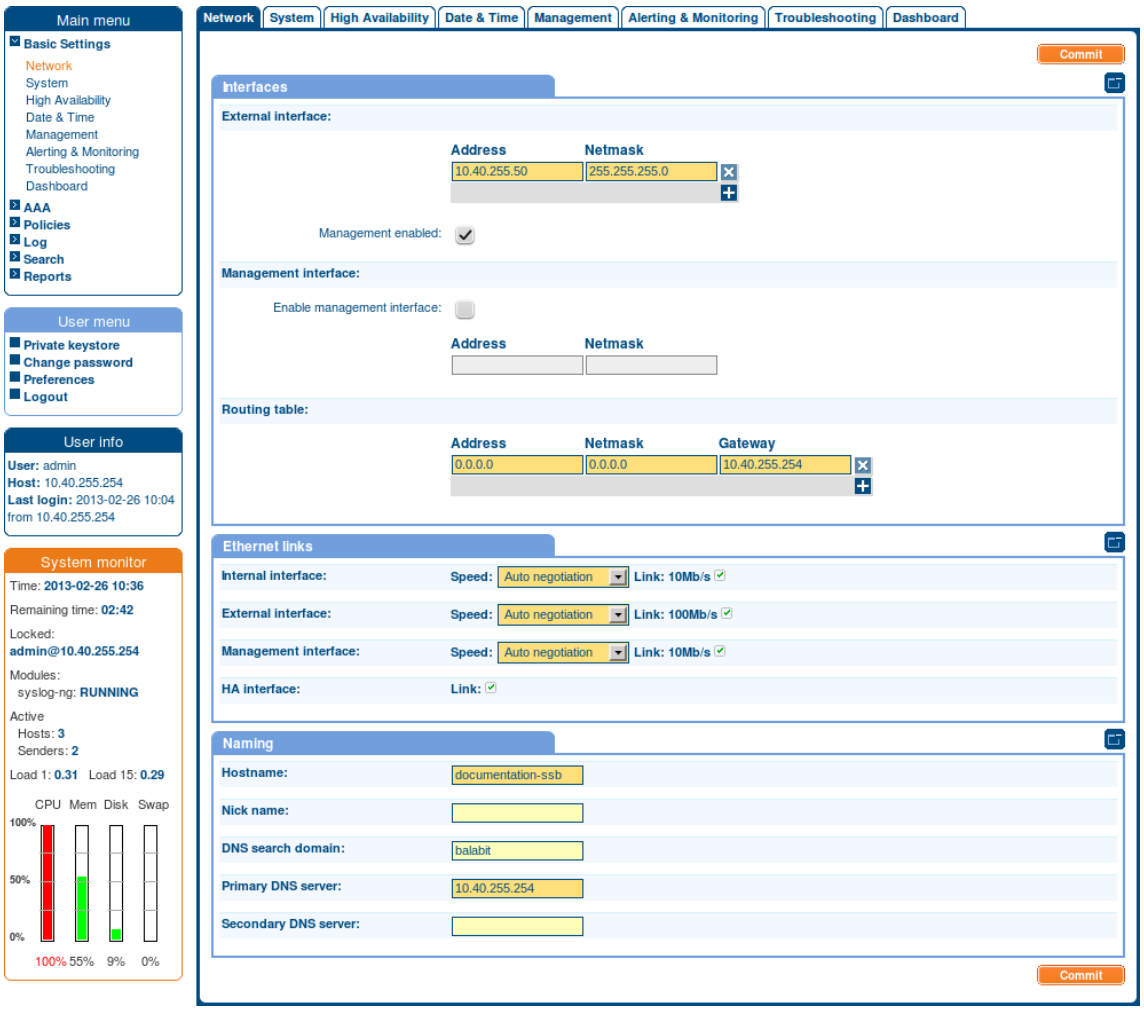

Figure 4.7. Network settings

- **External interface**: The Address and Netmask of the SSB network interface that receives client connections. Click the + and × icons to add new alias IP addresses (also called alias interfaces) or delete existing ones. At least one external interface must be configured. If the management interface is disabled, the SSB web interface can be accessed via the external interface. When multiple external interfaces are configured, the first one refers to the physical network interface, all others are alias interfaces. The SSB web interface can be accessed from all external interfaces (if no management interface is configured).

Optionally, you can enable access to the SSB web interface even if the management interface is configured by activating the **Management enabled** function.

Warning
If you enable management access on an interface and configure alias IP address(es) on the same interface, SSB will accept management connections only on the original address of the interface.

Note
Do not use IP addresses that fall into the following ranges:

- `1.2.0.0/16` (reserved for communication between SSB cluster nodes)
- `127.0.0.0/8` (localhost IP addresses)

Note
The speed of the interface is displayed for every interface. To explicitly set the speed of the interface, select the desired speed from the **Speed** field. Modifying the speed of the interface is recommended only for advanced users. Also note that changing the interface speed might not take effect if the network card of SSB has been replaced with one different from the original.

- **Management interface**: The Address and Netmask of the SSB network interface used to access the SSB web interface. If the management interface is configured, the web interface can be accessed only via this interface, unless access from other interfaces is explicitly enabled.

Note
Do not use IP addresses that fall into the following ranges:

- `1.2.0.0/16` (reserved for communication between SSB cluster nodes)
- `127.0.0.0/8` (localhost IP addresses)

4.3.1. Procedure – Configuring the management interface

Purpose:

To activate the interface, complete the following steps.

Steps:

Step 1. Navigate to **Basic Settings > Network > Interfaces**.

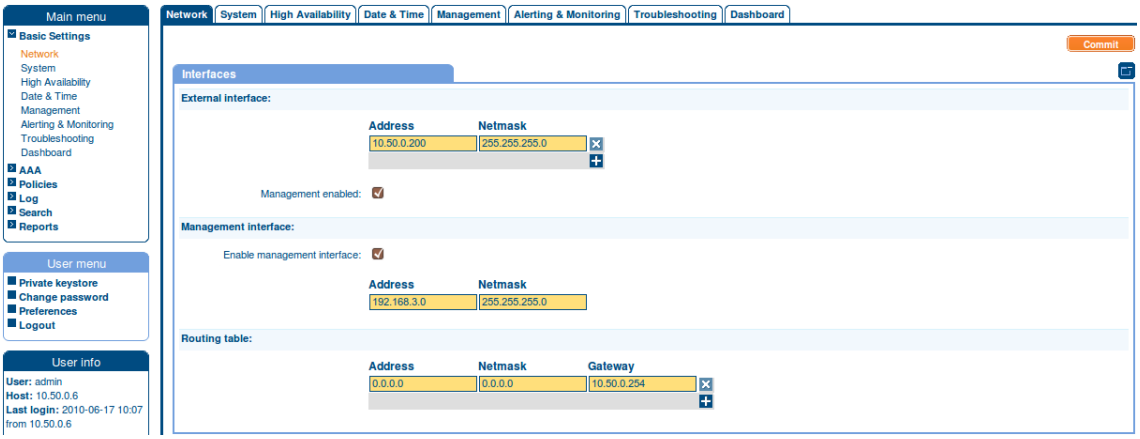

Figure 4.8. Configuring the management interface

Step 2. In the **Management interface** field, select **Enable management interface**.

Step 3. Into the **Address** field, enter the IP address of SSB's management interface.

Step 4. Into the **Netmask** field, enter the netmask related to the IP address.

Step 5.

Warning

After clicking Commit, the web interface will be available only via the management interface — it will not be accessible using the current (external) interface, unless the **Management enabled** option is selected for the external interface.

Ensure that the Ethernet cable is plugged and the management interface is connected to the network; this is indicated by a green check icon in the **Basic settings > Networks > Ethernet links > HA interface > Link** field. When using High Availability, ensure that the management interface of both SSB units is connected to the network.

The **HA interface** section indicates if a link is detected on the high availability interface.

Click Commit.

- **HA address**: The IP address of the high availability (HA) interface. Leave this field on *`Auto negotiation`* unless specifically requested by the support team.

Note

As of SSB version 1.1.1, when both nodes of a cluster boot up in parallel, the node with the `1.2.4.1` HA IP address will become the master node.

- **Interfaces > Routing table**: When sending a packet to a remote network, SSB consults the routing table to determine the path it should be sent. If there is no information in the routing table then the packet is sent to the default gateway. Use the routing table to define static routes to specific hosts or networks. You have to use the routing table if the internal interface is connected to multiple subnets, because the default gateway is (usually) towards the external interface. Click the ⊞ and ⊠ icons to

add new routes or delete existing ones. A route means that messages sent to the **Address/Netmask** network should be delivered to **Gateway**.

For detailed examples, see *Procedure 4.3.2, Routing management traffic to the management interface (p. 39)*.

- **Naming > Hostname**: Name of the machine running SSB.
- **Naming > Nick name**: The nickname of SSB. Use it to distinguish the devices. It is displayed in the core and boot login shells.
- **Naming > DNS search domain**: Name of the domain used on the network. When resolving the domain names of the audited connections, SSB will use this domain to resolve the target hostname if the appended domain entry of a target address is empty.
- **Naming > Primary DNS server**: IP address of the name server used for domain name resolution.
- **Naming > Secondary DNS server**: IP address of the name server used for domain name resolution if the primary server is unaccessible.

4.3.2. Procedure – Routing management traffic to the management interface

Purpose:

For security reasons — and also to reduce network usage on the external interface — it is recommended to direct all management-related traffic of SSB towards the management network interface. Such traffic includes access to the web interface, *backups* and *archiving*, *data forwarded to a remote destination* , and *e-mail or SNMP alerts* sent to the administrator.

Warning
Complete the following procedure only if the management interface is configured; otherwise the data sent by SSB will be lost. For details on configuring the management interface, see *Procedure 4.3.1, Configuring the management interface (p. 37)*.

Steps:

Step 1. To add a new routing entry, navigate to **Basic Settings > Network > Interfaces** and in the **Routing table** field, click .

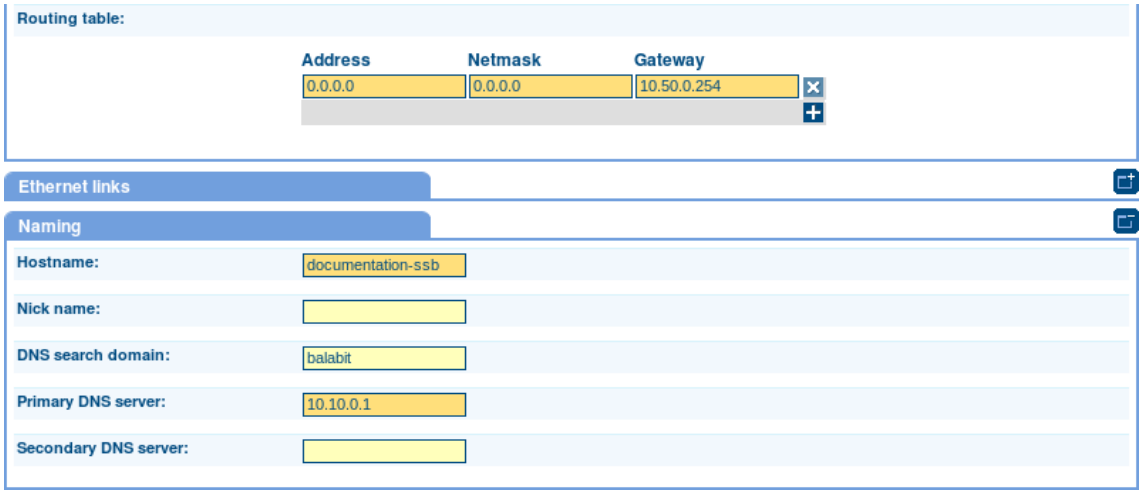

Figure 4.9. Routing

Step 2. Enter the IP address of the backup server (as set in *Procedure 4.7.1, Creating a backup policy (p. 55)*) into the **Address** field.

Step 3. Enter the related netmask into the **Netmask** field.

Step 4. Enter the IP address of the gateway used on that subnetwork into the **Gateway** field.

Step 5. Click Commit.

Step 6. Repeat Steps 1-5 and create a routing entry for other backup servers if needed.

Step 7. Repeat Steps 1-5 and create a routing entry for the SMTP server (as set in *Section 4.5, SNMP and e-mail alerts (p. 41)*).

Step 8. Repeat Steps 1-5 and create a routing entry for the remote destinations (as set in *Chapter 9, Forwarding messages from SSB (p. 157)*).

4.4. Date and time configuration

Date and time related settings of SSB can be configured on the **Date & Time** tab of the **Basic** page.

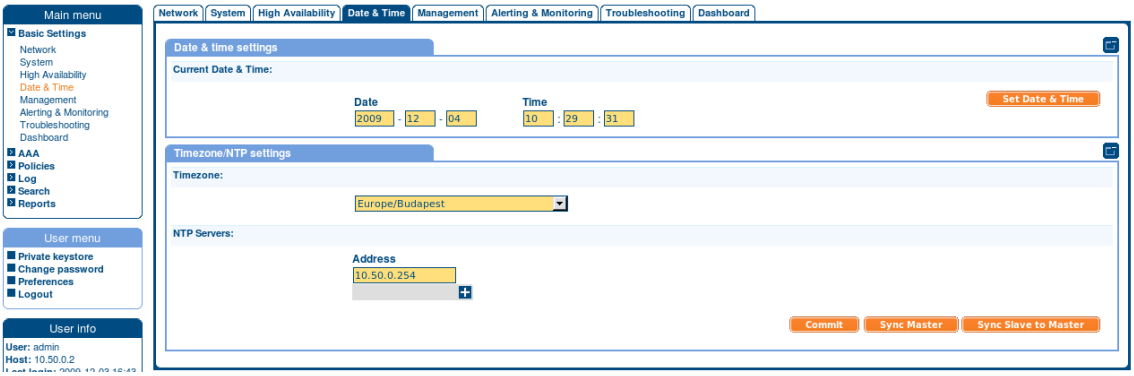

Figure 4.10. Date and time management

Warning
It is essential to set the date and time correctly on SSB, otherwise the date information of the logs will be inaccurate.

SSB displays a warning on this page and sends an alert if the time becomes out of sync.

To explicitly set the date and time on SSB, enter the current date into respective fields of the **Date & Time Settings** group and click **Set Date & Time**.

4.4.1. Procedure – Configuring a time (NTP) server

Purpose:

To retrieve the date automatically from a time server, complete the following steps.

Steps:

Step 1. Select your timezone in the **Timezone** field.

Step 2. Enter the IP address of an NTP time server into the **Address** field.

Step 3. Click Commit.

Step 4. Click the + and × icons to add new servers or delete existing ones.

Note
If the time setting of SSB is very inaccurate (that is, the difference between the system time and the actual time is great), it might take a long time to retrieve the date from the NTP server. In this case, click **Sync now** to sync the time immediately using SNTP.

When two SSB units are operating in high availability mode, the **Sync now** button is named **Sync Master**, and synchronizes the time of the master node to the NTP server. To synchronize the time between the master and the slave nodes, click **Sync Slave to Master**.

4.5. SNMP and e-mail alerts

E-mail alerts can be configured on the **Basic Settings > Management** page.

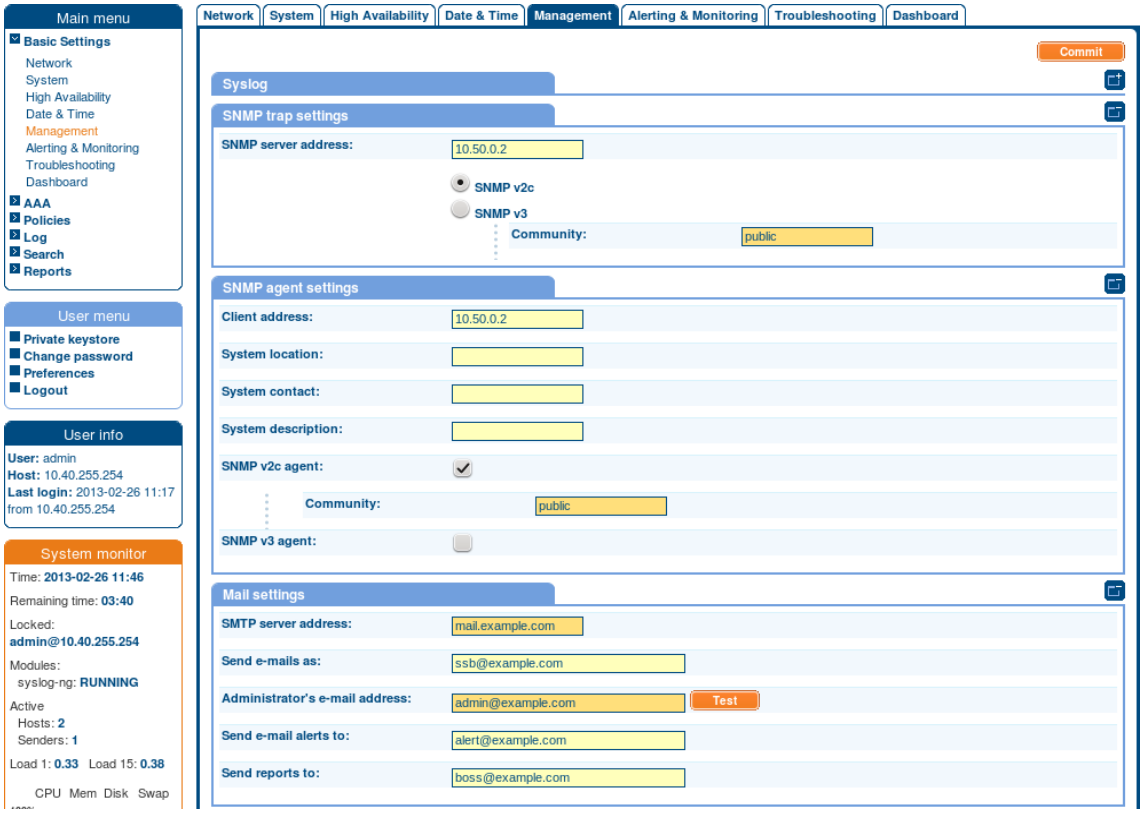

Figure 4.11. Configuring SNMP and e-mail alerts

4.5.1. Procedure – Configuring e-mail alerts

Purpose:

To configure e-mail alerts, complete the following steps:

Steps:

Step 1. Navigate to **Basic Settings > Management > Mail settings**.

Step 2. Enter the IP address or the hostname of the mail server into the **SMTP server address** field.

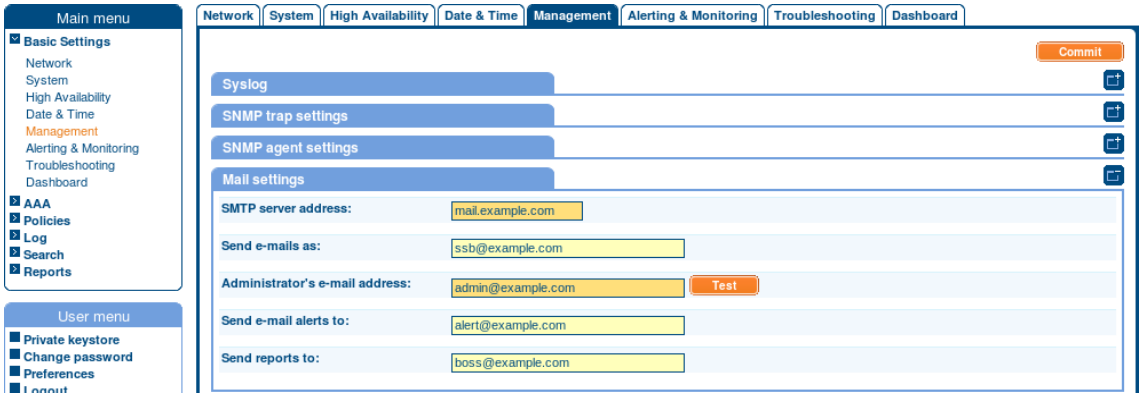

Figure 4.12. Configuring e-mail sending

Step 3. Enter the e-mail address where you want to receive e-mails from into the **Send e-mails as** field. This can be useful for e-mail filtering purposes. SSB sends e-mails from the address provided here. If no e-mail address is entered, e-mails will be sent from the default e-mail address.

Step 4. Enter the e-mail address of the administrator into the **Administrator's e-mail address** field. SSB sends notifications related to system-events (but not alerts and reports) to this address.

Step 5. Enter the e-mail address of the administrator into the **Send e-mail alerts to** field. SSB sends monitoring alerts to this address.

Step 6. Enter the e-mail address the person who should receive traffic reports from SSB into the **Send reports to** field. For details on reports, see *Section 13.7, Reports (p. 205)*.

Warning
To get alert e-mails, provide an e-mail address in this field. Sending alerts fails if these settings are incorrect, since the alerting e-mail address does not fall back to the administrator's e-mail address by default.

Step 7. Click Commit.

Step 8. Click **Test** to send a test message.

If the test message does not arrive to the server, check if SSB can access the server. For details, see *Chapter 16, Troubleshooting SSB (p. 222)*.

Step 9. Navigate to **Basic Settings** > **Alerting & Monitoring** and select in which situations should SSB send an e-mail alert. For details, see *Section 4.6, Configuring system monitoring on SSB (p. 47)*.

Step 10. Click Commit.

4.5.2. Procedure – Configuring SNMP alerts

Purpose:

SSB can send alerts to a central monitoring server via SNMP (Simple Network Management Protocol). To configure SNMP alerts, complete the following steps:

Steps:

Step 1. Navigate to **Basic Settings > Management > SNMP trap settings**.

Step 2. Enter the IP address or the hostname of the SNMP server into the **SNMP server address** field.

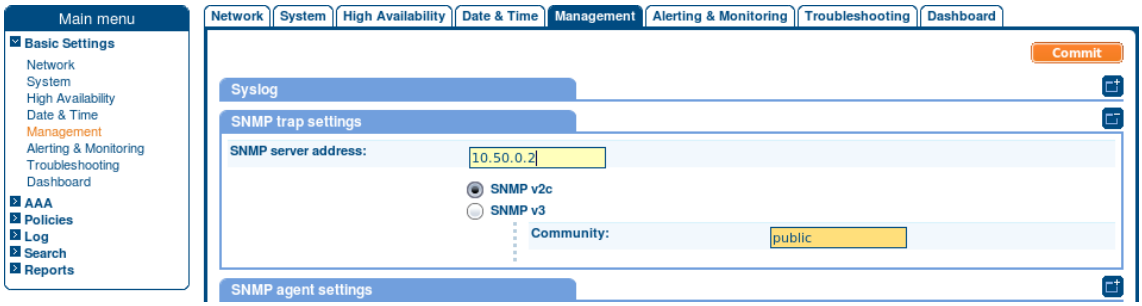

Figure 4.13. Configuring SNMP alerts

Step 3. Select the SNMP protocol to use.

- To use the SNMP v2c protocol for SNMP queries, select **SNMP v2c**, and enter the community to use into the **Community** field.

- To use the SNMP v3 protocol, select **SNMP v3** and complete the following steps:

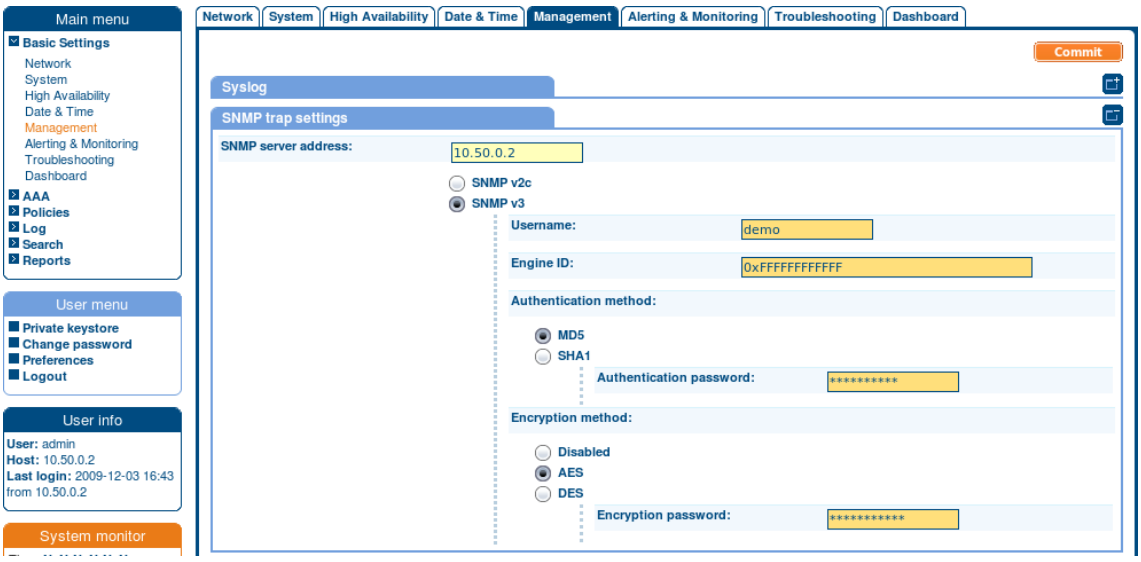

Figure 4.14. Configuring SNMP alerts using SNMPv3

Step a. Enter the username to use into the **Username** field.

Step b. Enter the engine ID to use into the **Engine ID** field. The engine ID is a hexadecimal number at least 10 digits long, starting with *0x*. For example *0xABABABABAB*.

Step c. Select the authentication method (**MD5 or SHA1**) to use from the **Authentication method** field.

Step d. Enter the password to use into the **Authentication password** field.

Step e. Select the encryption method (**Disabled, DES or AES**) to use from the **Encryption method** field.

Step f. Enter the encryption password to use into the **Encryption password** field.

Note
SSB accepts passwords that are not longer than 150 characters. The following special characters can be used:
`!"#$%&'()*+,-./:;<=>?@[\]^-`{|}`

Step 4. Click Commit.

Step 5. Navigate to **Basic Settings > Alerting & Monitoring** and select in which situations should SSB send an SNMP alert. For details, see *Section 4.6, Configuring system monitoring on SSB (p. 47)*.

Step 6. Click Commit.

4.5.3. Procedure – Querying SSB status information using agents

Purpose:

External SNMP agents can query the basic status information of SSB. To configure which clients can query this information, complete the following steps:

Steps:

Step 1. Navigate to **Basic Settings > Management > SNMP agent settings**.

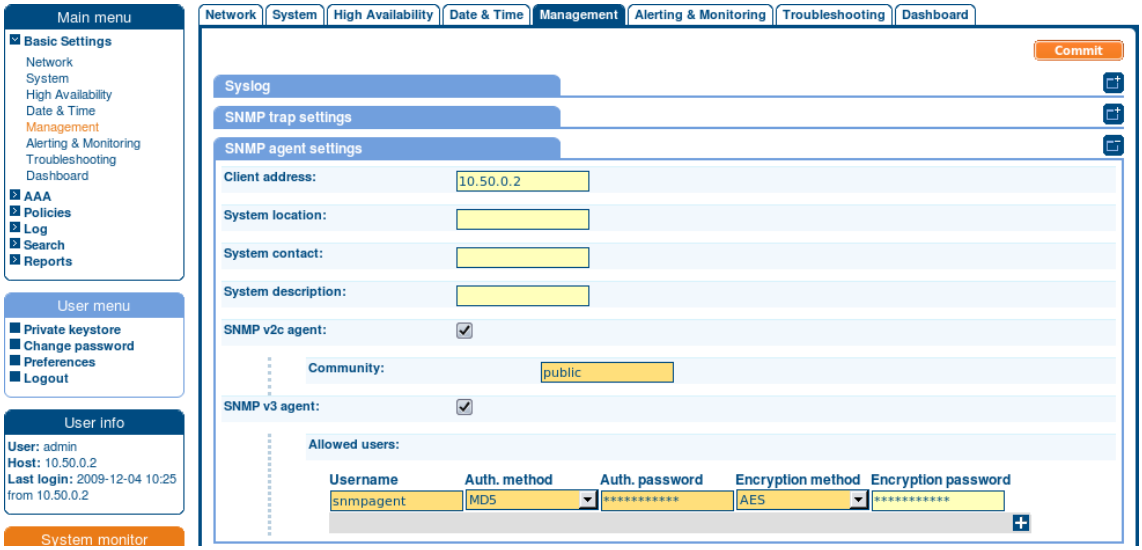

Figure 4.15. Configuring SNMP agent access

Step 2. The status of SSB can be queried dynamically via SNMP. By default, the status can be queried from any host. To restrict access to these data to a single host, enter the IP address of the host into the **Client address** field.

Step 3. Optionally, you can enter the details of the SNMP server into the **System location**, **System contact**, and **System description** fields.

Step 4. Select the SNMP protocol to use.

- To use the SNMP v2c protocol for SNMP queries, select **SNMP v2c agent**, and enter the community to use into the **Community** field.
- To use the SNMP v3 protocol, select **SNMP v3 agent** and complete the following steps:

Step a. Click ⊞

Step b. Enter the username used by the SNMP agent into the **Username** field.

Step c. Select the authentication method (**MD5 or SHA1**) to use from the **Auth. method** field.

Step d. Enter the password used by the SNMP agent into the **Auth. password** field.

Step e. Select the encryption method (**Disabled, DES or AES**) to use from the **Encryption method** field.

Step f. Enter the encryption password to use into the **Encryption password** field.

Step g. To add other agents, click ⊞.

Note
SSB accepts passwords that are not longer than 150 characters. The following special characters can be used:
`!"#$%&'()*+,-./:;<=>?@[\]^-`{|}`

Step 5. Click Commit.

4.6. Configuring system monitoring on SSB

SSB continuously monitors a number of parameters of the SSB hardware and its environment. If a parameter reaches a critical level (set in its respective **Maximum** field), SSB sends e-mail and SNMP messages to alert the administrator.

SSB sends SNMP alerts using the management network interface by default, or using the external interface if the management interface is disabled. SSB supports the SNMPv2c and SNMPv3 protocols. The SNMP server set on the **Management** tab can query status information from SSB.

Tip
To have your central monitoring system recognize the SNMP alerts sent by SSB, select **Basic Settings** > **Alerting & Monitoring** > **Download MIBs** to download the SSB-specific Management Information Base (MIB), then import it into your monitoring system.

Figure 4.16. Configuring SNMP and e-mail alerting

4.6.1. Procedure – Configuring monitoring

Purpose:

To configure monitoring, complete the following steps:

Steps:

Step 1. Navigate to **Basic Settings > Alerting & Monitoring**.

Step 2. The default threshold values of the parameters are suitable for most situations. Adjust the thresholds only if needed.

Step 3. Click Commit.

Step 4. Navigate to **Basic Settings > Management** and verify that the **SNMP settings** and **Mail settings** of SSB are correct. SSB sends alerts only to the alert e-mail address and to the SNMP server.

Warning
Sending alerts fails if these settings are incorrect.

The following sections describe the parameters you can receive alerts on.

- For details on health-monitoring alerts, see *Section 4.6.2, Health monitoring (p. 49)*.
- For details on system-monitoring alerts, see *Section 4.6.5, System related traps (p. 52)*.
- For details on syslog-related alerts, see *Section 4.6.6, Alerts related to syslog-ng (p. 54)*.

4.6.2. Health monitoring

- **Disk utilization maximum**: Ratio of free space available on the hard disk. SSB sends an alert if the log files use more space than the set value. Archive the log files to a backup server to free disk space. For details, see *Section 4.8, Archiving and cleanup (p. 59)*.

Note
The alert message includes the actual disk usage, not the limit set on the web interface. For example, you set SSB to alert if the disk usage increases above *10* percent. If the disk usage of SSB increases above this limit (for example to *17* percent), you receive the following alert message: *less than 90% free (= 17%)*. This means that the amount of used disk space increased above 10% (what you set as a limit, so it is less than 90%), namely to 17%.

- **Load 1|5|15 maximum**: The average load of SSB during the last one, five, or 15 minutes.
- **Swap utilization maximum**: Ratio of the swap space used by SSB. SSB sends an alert if it uses more swap space than the set value.

4.6.3. Procedure – Preventing disk space fill up

Purpose:

To prevent disk space from filling up, complete the following steps:

Steps:

Step 1. Navigate to **Basic Settings > Management > Disk space fill up prevention**.

Step 2. Set the limit of maximum disk utilization in percents in the respective field. When disk space is used above the set limit, SSB disconnects all clients. Entering *0* turns the feature off. The default value is *0*.

Step 3. *Optional step*: Enable the **Automatically start archiving** option to automatically start all configured archiving/cleanup jobs when disk usage goes over the limit.

Note
If there is no archiving policy set, enabling this option will not trigger automatic archiving.

Step 4. Navigate to **Basic Settings > Alerting & Monitoring > System related traps** and enable alert **Disk usage is above the defined ratio**.

Step 5. Click Commit.

4.6.4. Procedure – Configuring message rate alerting

Purpose:

With message rate alerting, you can detect the following abnormalities in SSB:

- The syslog-ng inside SSB has stopped working.
- One of the clients/sites sending logs is not detectable.
- One of the clients/sites is sending too many logs, probably unnecessarily.

Message rate alerting can be set for sources, spaces and destinations (remote or local).

Steps:

Step 1. Navigate to **Log** and select **Sources**, **Spaces** or **Destinations**.

Step 2. Enable **Message rate alerting**.

Step 3. In case of **Sources**, select the counter to be measured:

- *Messages*: Number of messages
- *Messages/sender*: Number of messages per sender (the last hop)
- *Messages/hostname*: Number of messages per host (based on the hostname in the message)

In case of **Spaces** or **Destinations**, the counter is the number of messages.

Step 4. Select the time period (between 5 minutes and 24 hours) during which the range is to be measured.

Step 5. Enter the range that is considered normal in the **Minimum** and **Maximum** fields.

Step 6. Select the alerting frequency in the **Alert** field. **Once** sends only one alert (and after the problem is fixed, a "Fixed" message), **Always** sends an alert each time the result of the measurement falls outside the preset range.

Example 4.2. Creating an early time alert
In case you want an early time alert, can create a normal (non master) alert with a very low minimum number of messages and a low check interval.

Counter	Period	Minimum	Maximum	Alert	Master alert
Messages	24 hours	10000	1000000	Once	☐
Messages	30 minutes	10	1000	Once	☐

Figure 4.17. Creating an early time alert

Step 7. If you have set more than one message rate alerts, you can set a master alert where applicable. To set an alert to be a master alert, select the **Master alert** checkbox next to it.

When a master alert is triggered (and while it remains triggered), all other alerts for the given source/destination/space are suppressed. A master alert only blocks the other alerts that would be triggered at the given timeslot. A 24-hour alert does not block alerts that would be triggered at, for example 00:05.

Suggestions for setting the master alert:

- set the master alert to low check interval (5 minutes, if possible)
- set the master alert to a lower check interval than the alerts it supresses
- set the master alert to have more lax limits than the alerts it supresses

The following examples demonstrate a few common use cases of a **Master alert**.

Example 4.3. Using the master alert to indicate unexpected events
The user has 2 relays (sender) and 10 hosts per each relay (=20 hosts). Each host sends approximately 5-10 messages in 5 minutes. Two message rate alerts are set, and one master alert to signal extreme unexpected events. Such event can be that either a host is undetectable and probably has stopped working, or that it sends too many logs, probably due to an error. The following configuration helps detecting these errors without having to receive hundreds of alerts unnecessarily.

Counter	Period	Minimum	Maximum	Alert	Master alert
Messages/hostnar	5 minutes	50	100	Once	☐
Messages/sender	5 minutes	500	1000	Once	☐
Messages	5 minutes	0	10000	Once	☑

Figure 4.18. Using the master alert to indicate unexpected events

Step 8. *Optional step*: Global alerts count the number of all messages received by syslog-ng on all sources, including internal messages.

Step a. Navigate to **Log > Options > Message rate alerting statistics**. To add a global alert, click ⊕ at **Global alerts**.

Step b. Select the time period (between 5 minutes and 24 hours) during which the range is to be measured.

Step c. Enter the range that is considered normal in the **Minimum** and **Maximum** fields.

Step d. Select the alerting frequency in the **Alert** field. **Once** sends only one alert (and after the problem is fixed, a "Fixed" message), **Always** sends an alert each time the result of the measurement falls outside the preset range.

Step e. To set the alert as a system-wide master alert, select **Global master alert**. It will suppress all other log rate alerts on SSB when it is triggered.

Note
In the following cases, a so-called "always"-type super-master alert is triggered automatically.

If all or some of the statistics from syslog-ng cannot be fetched, an alert is sent out and all other errors are suppressed until the error is fixed.

If, for some reason, syslog-ng sends an unprocessable amount of statistics (for example because of some invalid input data), a similar super-master alert is triggered and stops processing the input.

Step 9. *Optional step*: Navigate to **Log > Options > Message rate alerting statistics**. Set the maximum number of alerts you want to receive in **Limit of alerts sent out in a batch** to prevent alert flooding. SSB will send alerts up to the predefined value and then one single alert stating that too many message alerts were generated and the excess amount have not been sent.

Warning
Hazard of data loss! The alerts over the predefined limit will be unreachable.

4.6.5. System related traps

Name	SNMP alert ID	Description
Login failed	*xcbLoginFailure*	Failed login attempts from SSB web interface.
Successful login	*xcbLogin*	Successful login attempts into SSB web interface.
Logout from the management interface	*xcbLogout*	Logouts from SSB web interface.
Configuration changed	*xcbConfigChange*	Any modification of SSB's configuration.

Name	SNMP alert ID	Description
General alert	`xcbAlert`	General alerts and error messages occurring on SSB. Note, that alerts on general alerts and errors are sent whenever there is an alert or error level message in the SSB system log. These messages are very verbose and mainly useful only for debugging purposes. Enabling these alerts may result in multiple e-mails or SNMP traps sent about the same event.
General error	`xcbError`	
Data and configuration backup failed	`xcbBackupFailed`	Alerts if the backup procedure is unsuccessful.
Data archiving failed	`xcbArchiveFailed`	Alerts if the archiving procedure is unsuccessful.
Database error occurred	`xcbDBError`	An error occurred in the database where SSB stores alerts and accounting information. Contact our support team (see *Section 5, Contact and support information (p. xiii)* for contact information).
License limit reached	`xcbLimitReached`	Maximum number of clients has been reached.
HA node state changed	`xcbHaNodeChanged`	A node of the SSB cluster changed its state, for example, a takeover occurred.
Timestamping error occured	`xcbTimestampError`	An error occurred during the timestaming process, for example the timestamping server did not respond.
Time sync lost	`xcbTimeSyncLost`	The system time became out of sync.
Raid status changed	`xcbRaidStatus`	The status of the node's RAID device changed its state.
Hardware error occured	`xcbHWError`	SSB detected a hardware error.
Firmware is tainted	`xcbFirmwareTainted`	A user has locally modified a file from the console.

Name	SNMP alert ID	Description
Disk usage is above the defined ratio	*xcbDiskFull*	Disk space is used above the limit set in **Disk space fill up prevention**.

Table 4.1. System related traps

4.6.6. Alerts related to syslog-ng

Name	SNMP alert ID	Description
syslog-ng failure	*syslogngFailureTrap*	The syslog-ng application did not start properly, shut down unexpectedly, or encountered another problem. Depending on the error, SSB may not accept incoming messages or send them to the destinations.
Remote syslog-ng peer configuration changed	*peerConfigChangeTrap*	The configuration of the syslog-ng application running on a remote host that sents its logs to SSB has been changed. Note that such changes are detected only if the remote peer uses at least version 3.0 of syslog-ng or version 3.0 of the syslog-ng Agent, and if messages from the *internal* source are sent to SSB.
Logspace exceeded warning size	*spaceSizeLimit*	The size of a log space has exceeded the size set as warning limit.
Message rate was outside the specified limits	*ssbAbsoluteMessageRateAlert*	The message rate has exceeded the minimum or maximum value.
Too many message rate alerts were generated	*ssbRateLimitTooManyAlerts*	SSB is generating too many message rate alerts, probably due to unusual traffic that may need investigation and further user actions.
Error during syslog-ng traffic statistics processing	*ssbStatisticsError*	There was an error during querying and processing statistics of incoming, forwarded, stored, and dropped messages.
Error during an sql-source related operation	*ssbSqlSourceAlert*	It is not possible to connect or log in to the SQL server, the SQL table is not found, or there is a problem with executing SQL queries, for example

Name	SNMP alert ID	Description
		insufficient permissions to access the database.

Table 4.2. Alerts related to syslog-ng

4.7. Data and configuration backups

Backups create a snapshot of SSB's configuration or the data which can be used for recovery in case of errors. The syslog-ng Store Box can create automatic backups of its configuration and the stored logs to a remote server. To configure backups, you first have to create a backup policy, and then set it as a system backup policy (for configuration backups) or data backup policy (for logspace backups).

Backup deletes all other data from the target directory; restoring a backup deletes all other data from SSB. For details on restoring configuration and data from backup, see *Procedure 16.7, Restoring SSB configuration and data (p. 235)*.

4.7.1. Procedure – Creating a backup policy

Backup policies define the address of the backup server, which protocol to use to access it, and other parameters.

Steps:

Step 1. Navigate to **Policies** > **Backup & Archive/Cleanup** and click ➕ in the **Backup policies** section to create a new backup policy.

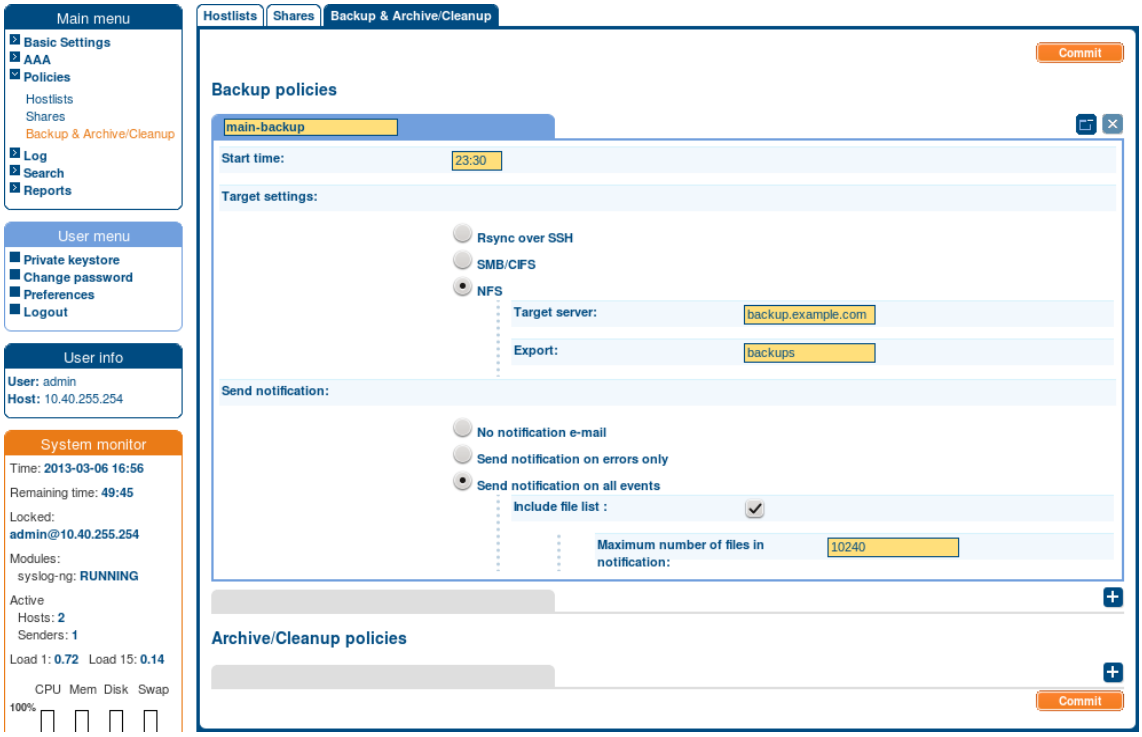

Figure 4.19. Configuring backups

Step 2. Enter a name for the backup policy (for example *config-backup*).

Step 3. Enter the time when the backup process should start into the **Start time** field in HH:MM format (for example *23:00*).

Step 4. Enter the IP address or the hostname of the remote server into the **Target server** field (for example *backup.example.com*).

Step 5. SSB can access the remote server via different protocols. Select the one to use from the available protocols:

- **Rsync over SSH**: Execute the `rsync` command via the Secure Shell protocol. Note that the backup server must run rsync version 3.0 or newer.

- **SMB/CIFS**: Server Message Block protocol used on Microsoft Windows Network.

Warning
The CIFS implementation of NetApp storage devices is not compatible with the CIFS implementation used in SSB, therefore it is not possible to create backups and archives from SSB to NetApp devices using the CIFS protocol (the operation fails with a similar error message: *`/opt/scb/mnt/14719217504d41370514043/reports/2010": Permission denied (13) '2010/day/' rsync: failed to set times on`*).

To overcome this problem, either:

- use the NFS protocol to access your NetApp devices, or
- use a backup device that has a CIFS implementation compatible with SSB, for example, Windows or Linux Samba.

Warning
When using the CIFS protocol to backup or archive files to a target server running Windows 2008 R2 that uses NTLMv2 authentication, the operation may fail with a similar error message:

```
CIFS VFS: Unexpected SMB signature
Status code returned 0xc000000d NT_STATUS_INVALID_PARAMETER
CIFS VFS: Send error in SessSetup = -22
CIFS VFS: cifs_mount failed w/return code = -22
CIFS VFS: Server requires packet signing to be enabled in
/proc/fs/cifs/SecurityFlags.
CIFS VFS: cifs_mount failed w/return code = -95
CIFS VFS: Server requires packet signing to be enabled in
/proc/fs/cifs/SecurityFlags.
CIFS VFS: cifs_mount failed w/return code = -95
```

To overcome this problem, either:

- use the NFS protocol to access your Windows 2008 R2 servers, or
- edit the registry of the Windows 2008 R2 server or apply a hotfix. For details, see *Article 957441* in the Microsoft® Support site.

- **NFS**: Network File System protocol.

Step 6. Provide the protocol-specific parameters for the selected method. The protocol-specific parameters are described in *Section 4.9, Parameters of the backup and archiving protocols (p. 62)*.

Step 7. To receive e-mail notification of the backup, select the **Send notification on errors only** or the **Send notification on all events** option. Notifications are sent to the administrator e-mail address set on the **Management** tab.

To include the list of files in the e-mail, select **Send notification on all events** and enable the **Include file list** option. However, note that if list is very long, the SSB web interface might become unaccessible. In this case, set the **Maximum number of files in notification** lower. After this number has been reached, file names will be omitted from the notification.

Note
This e-mail notification is different from the one set on the **Alerting & Monitoring** tab. This notification is sent to the administrator's e-mail address, while the alerts are sent to the alert e-mail address (see *Section 4.6, Configuring system monitoring on SSB (p. 47)*).

Step 8. Click Commit.

4.7.2. Procedure – Creating configuration backups

To create a configuration backup, assign a backup policy as the **System backup policy** of SSB.

Tip
To create an immediate backup of SSB's configuration to your machine (not to the backup server), select **Basic Settings > System > Export configuration**. Note that the configuration export contains only the system settings and configuration files (including changelogs). System backups includes additional information like reports and alerts.

Steps:

Step 1. Navigate to **Basic Settings > Management > System backup**.

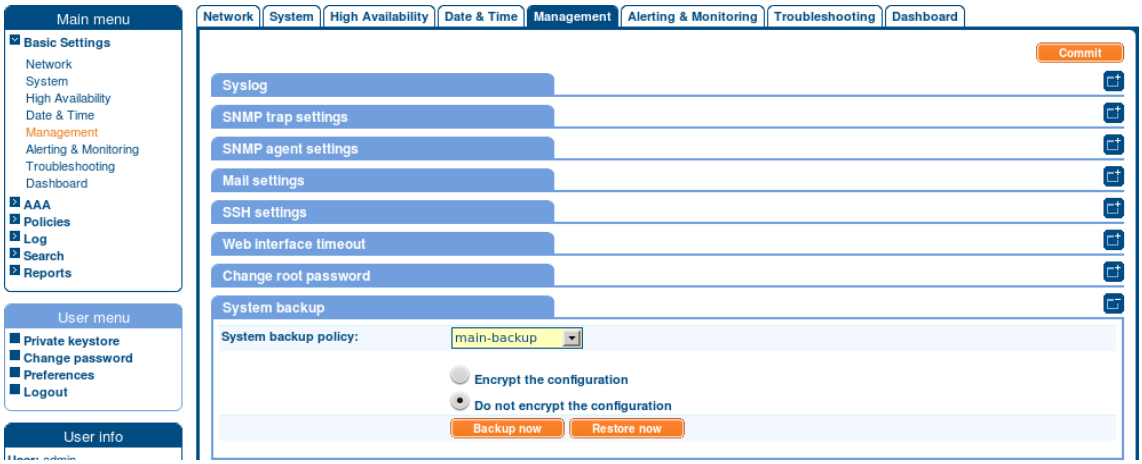

Figure 4.20. Configuring system backups

Step 2. Select the backup policy you want to use for backing up the configuration of SSB in the **System backup policy** field.

Step 3. Click Commit.

Step 4. *Optional:* To start the backup process immediately, click **Backup now**. The **Backup now** functionality works only after a backup policy has been selected and committed.

4.7.3. Procedure – Creating data backups

To configure data backups, assign a backup policy to the logspace.

Steps:

Step 1. Navigate to **Log > Spaces**.

Step 2. Select a backup policy in the **Backup policy** field.

Step 3. Click Commit.

Step 4. *Optional:* To start the backup process immediately, click **Backup** or **Backup ALL**. The **Backup** and **Backup ALL** functionalities work only after a backup policy has been selected and committed.

4.7.4. Procedure – Encrypting configuration backups with GPG

You can encrypt the configuration file of SSB during system backups using the public-part of a GPG key. The system backups of SSB contain other information as well (for example, databases), but only the configuration file is encrypted. Note that system backups do not contain logspace data.

The GPG key you upload must be permitted to encrypt data. Keys that can be used only for signing cannot be used to encrypt the configuration file.

Note
It is not possible to directly import a GPG-encrypted configuration into SSB, it has to be decrypted locally first.

For details on restoring configuration from a configuration backup, see *Procedure 16.7, Restoring SSB configuration and data (p. 235)*

Steps:

Step 1. Navigate to **Basic > System > Management > System backup**.

Step 2. Select **Encrypt configuration**.

Step 3. Select .

- To upload a key file, click **Browse**, select the file containing the public GPG key, and click **Upload**. SSB accepts both binary and ASCII-armored GPG keys.
- To copy-paste the key from the clipboard, paste it into the **Key** field and click **Set**.

Step 4. Click Commit.

4.8. Archiving and cleanup

Archiving transfers data from SSB to an external storage solution, cleanup removes (deletes) old files. Archived data can be accessed and searched, but cannot be restored (moved back) to the SSB appliance.

To archive data, you have to create an archive policy and assign it to the logspace you want to archive.

4.8.1. Procedure – Creating an archive policy

Archive policies defines the address of the remote storage, which protocol to use to access it, and other parameters.

Step 1. Navigate to **Policies > Backup & Archive/Cleanup** and click in the **Archive/Cleanup policies** section to create a new archive policy.

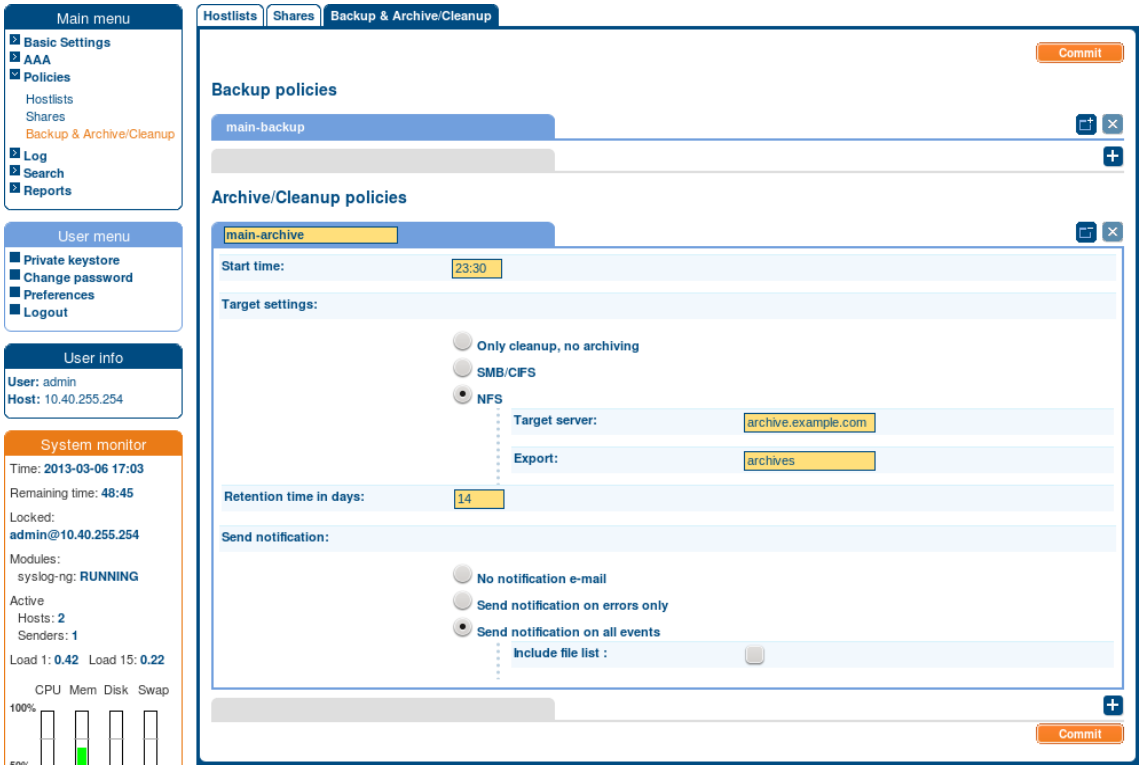

Figure 4.21. Configuring backups and archiving

Step 2. Enter a name for the archive policy.

Step 3. Enter the time when the archive process should start into the **Start time** field in HH:MM format (for example *23:00*).

Step 4. SSB can access the remote server via different protocols. Select which one to use:

- **Only cleanup, no archiving**: Do not archive data to a server, simply delete the data that is older than **Retention time in days**.

Warning
No archiving permanently deletes all log files and data that is older than **Retention time in days** without creating a backup copy or an archive. Such data is irrecoverably lost. Use this option with care.

This setting does not delete existing archives from an external CIFS or NFS server.

- **SMB/CIFS**: Server Message Block protocol used on Microsoft Windows Network.

Warning
The CIFS implementation of NetApp storage devices is not compatible with the CIFS implementation used in SSB, therefore it is not possible to create backups and archives from SSB to NetApp devices using the CIFS protocol (the operation fails with a similar error message: *`/opt/scb/mnt/14719217504d41370514043/reports/2010": Permission denied (13) '2010/day/' rsync: failed to set times on`*).

To overcome this problem, either:

- use the NFS protocol to access your NetApp devices, or
- use a backup device that has a CIFS implementation compatible with SSB, for example, Windows or Linux Samba.

Warning
When using the CIFS protocol to backup or archive files to a target server running Windows 2008 R2 that uses NTLMv2 authentication, the operation may fail with a similar error message:

```
CIFS VFS: Unexpected SMB signature
Status code returned 0xc000000d NT_STATUS_INVALID_PARAMETER
CIFS VFS: Send error in SessSetup = -22
CIFS VFS: cifs_mount failed w/return code = -22
CIFS VFS: Server requires packet signing to be enabled in
/proc/fs/cifs/SecurityFlags.
CIFS VFS: cifs_mount failed w/return code = -95
CIFS VFS: Server requires packet signing to be enabled in
/proc/fs/cifs/SecurityFlags.
CIFS VFS: cifs_mount failed w/return code = -95
```

To overcome this problem, either:

- use the NFS protocol to access your Windows 2008 R2 servers, or
- edit the registry of the Windows 2008 R2 server or apply a hotfix. For details, see *Article 957441* in the Microsoft® Support site.

- **NFS**: Network File System protocol.

Warning
If you modify the connection protocol of an existing policy (for example, from NFS to SMB/CIFS), the old archives will become inaccessible. To avoid this, create a new archive policy instead, using the new connection protocol, and configure it for all affected logspaces (**Log** > **Spaces** > **<name-of-the-logspace>** > **Archive/Cleanup policy**). This way, both the old and the new archived logs will be accessible.

Step 5. Enter the IP address or the hostname of the remote server into the **Target settings** > **Target server** field (for example `backup.example.com`).

Step 6. Fill the **Retention time in days** field. Data older than this value is archived to the external server.

Note
The archived data is deleted from SSB.

Step 7. Provide the protocol-specific parameters for the selected method. The protocol-specific parameters are described in *Section 4.9, Parameters of the backup and archiving protocols (p. 62)*.

Step 8. To receive e-mail notifications, select the **Send notification on errors only** or the **Send notification on all events** option. Notifications are sent to the administrator e-mail address set on the **Management** tab, and include the list of the files that were backed up.

Note
This e-mail notification is different from the one set on the **Alerting & Monitoring** tab. This notification is sent to the administrator's e-mail address, while the alerts are sent to the alert e-mail address (see *Section 4.6, Configuring system monitoring on SSB (p. 47)*).

Step 9. Click Commit.

4.8.2. Procedure – Archiving the collected data

To configure data archiving, assign an archive policy to the logspace you want to archive.

Steps:

Step 1. Navigate to **Log > Spaces**.

Step 2. Select the logspace you want to archive.

Step 3. Select the archive policy you want to use in the **Archive/Cleanup policy** field.

Step 4. Click Commit.

Step 5. *Optional:* To start the archiving process immediately, click **Archive now**. The **Archive now** functionality works only after the archiving has been configured.

4.9. Parameters of the backup and archiving protocols

This section describes the details of the protocols used for data backup and archiving.

- For details on using Rsync, see *Procedure 4.9.1, Configuring Rsync over SSH (p. 62)*.
- For details on using the Samba protocol, see *Procedure 4.9.2, Configuring SMB (p. 64)*.
- For details on using NFS, see *Procedure 4.9.3, Configuring NFS (p. 66)*.

4.9.1. Procedure – Configuring Rsync over SSH

Purpose:

The **Rsync over SSH** backup method connects the target server with SSH and executes the `rsync` UNIX command to copy the data to the remote server. SSB authenticates itself with a public key — password-based authentication is not supported. To configure this method, complete the following steps.

Warning
The backup server must run rsync version 3.0 or newer.

Steps:

Step 1. Select **Rsync over SSH** from the **Target settings** radio buttons.

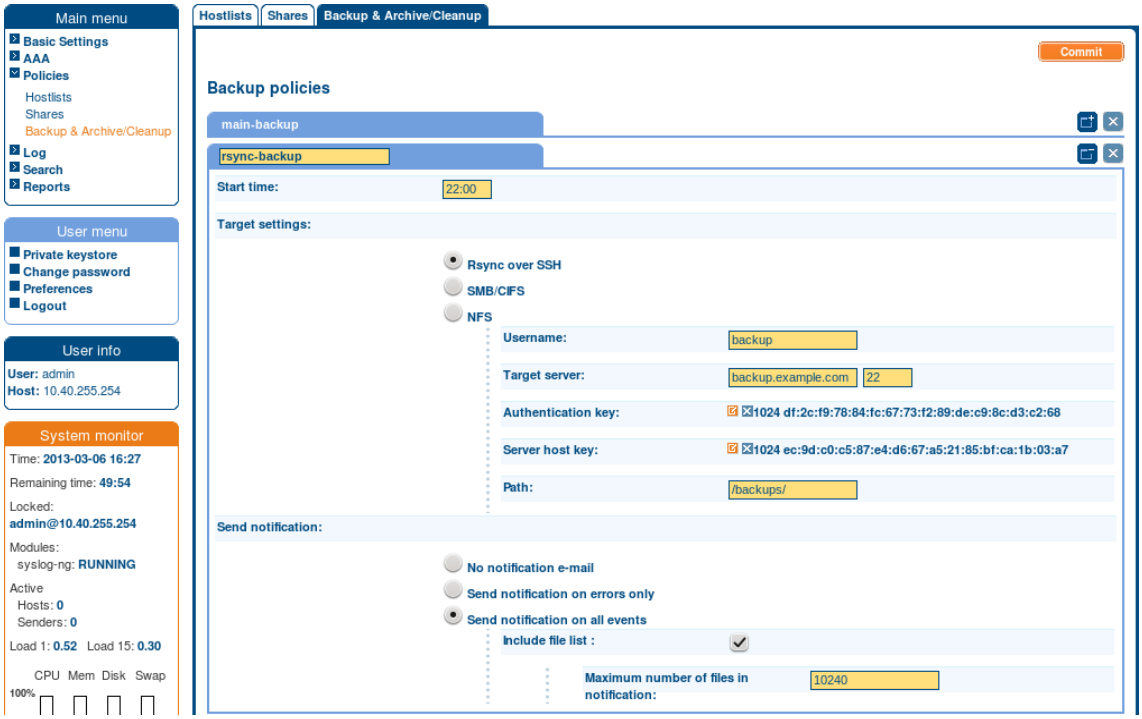

Figure 4.22. Configuring backups using rsync

Step 2. Enter the username used to logon to the remote server into the **Username** field.

Step 3. Click in the **Authentication key** field. A popup window is displayed.

Step 4. Generate a new keypair by clicking **Generate** or upload or paste an existing one. This key will be used to authenticate SSB on the remote server. The public key of this keypair must be imported to the remote server.

Step 5. Click in the **Server host key** field. A popup window is displayed.

Step 6. Click **Query** to download the host key of the server, or upload or paste the host key manually. SSB will compare the host key shown by the server to this key, and connect only if the two keys are identical.

Figure 4.23. Configuring SSH keys

Step 7. Enter the port number of the SSH server running on the remote machine into the **Port** field.

Step 8. Enter the path to the backup directory on the target server into the **Path** field (for example `/backups`). SSB saves all data into this directory, automatically creating subdirectories for logspaces. As a result of this, the same backup policy can be used for multiple logspaces. To ensure that a restore can be performed even if the logspace has been renamed, the subdirectories are created using a persistent internal ID of the logspace. To facilitate manual debugging, a text file is also saved in the directory with the name of the logspace, containing the internal ID for the logspace. This text file is only provided for troubleshooting purposes and is not used by SSB in any way.

Step 9. Click Commit.

4.9.2. Procedure – Configuring SMB

Purpose:

The **SMB/CIFS** backup method connects to a share on the target server with Server Message Block protocol. SMB/CIFS is mainly used on Microsoft Windows Networks. To configure this method, complete the following steps.

Warning
The CIFS implementation of NetApp storage devices is not compatible with the CIFS implementation used in SSB, therefore it is not possible to create backups and archives from SSB to NetApp devices using the CIFS protocol (the operation fails with a similar error message: `/opt/scb/mnt/14719217504d41370514043/reports/2010": Permission denied (13) '2010/day/' rsync: failed to set times on`).

To overcome this problem, either:

- use the NFS protocol to access your NetApp devices, or
- use a backup device that has a CIFS implementation compatible with SSB, for example, Windows or Linux Samba.

Warning
When using the CIFS protocol to backup or archive files to a target server running Windows 2008 R2 that uses NTLMv2 authentication, the operation may fail with a similar error message:

```
CIFS VFS: Unexpected SMB signature
Status code returned 0xc000000d NT_STATUS_INVALID_PARAMETER
CIFS VFS: Send error in SessSetup = -22
CIFS VFS: cifs_mount failed w/return code = -22
CIFS VFS: Server requires packet signing to be enabled in /proc/fs/cifs/SecurityFlags.
CIFS VFS: cifs_mount failed w/return code = -95
CIFS VFS: Server requires packet signing to be enabled in /proc/fs/cifs/SecurityFlags.
CIFS VFS: cifs_mount failed w/return code = -95
```

To overcome this problem, either:

- use the NFS protocol to access your Windows 2008 R2 servers, or
- edit the registry of the Windows 2008 R2 server or apply a hotfix. For details, see *Article 957441* in the Microsoft® Support site.

Steps:

Step 1. Select **SMB/CIFS** from the **Target settings** radio buttons.

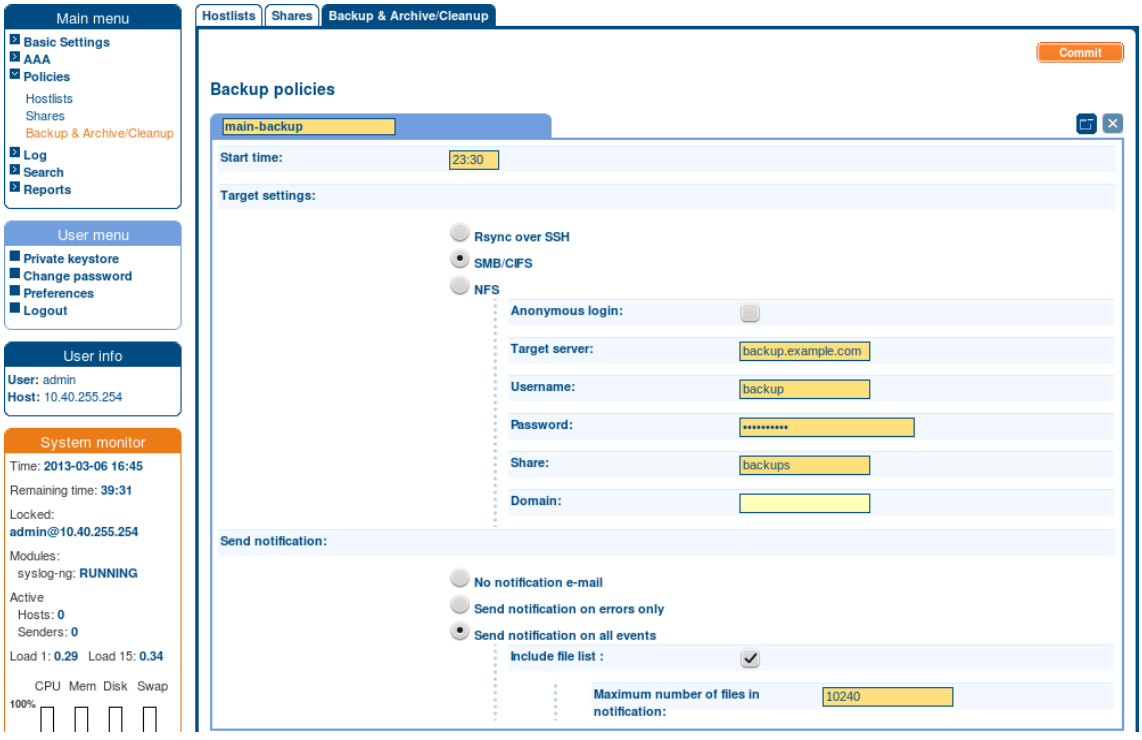

Figure 4.24. Configuring backups via SMB/CIFS

Step 2. Enter the username used to logon to the remote server into the **Username** field, or select the **Anonymous** option.

Step 3. Enter the password corresponding to the username into the **Password** field.

Note
SSB accepts passwords that are not longer than 150 characters. The following special characters can be used:
`!"#$%&'()*+,-./:;<=>?@[\]^-`{|}`

Step 4. Enter the name of the share into the **Share** field.

SSB saves all data into this directory, automatically creating the subdirectories. Backups of log files are stored in the *data*, configuration backups in the *config* subdirectory.

Step 5. Enter the domain name of the target server into the **Domain** field.

Step 6. Click Commit.

4.9.3. Procedure – Configuring NFS

Purpose:

The **NFS** backup method connects to a shared directory of the target server with the Network File Share protocol. To configure this method, enter the name of the NFS export into the **Export** field. SSB saves all data into this directory, automatically creating the subdirectories.

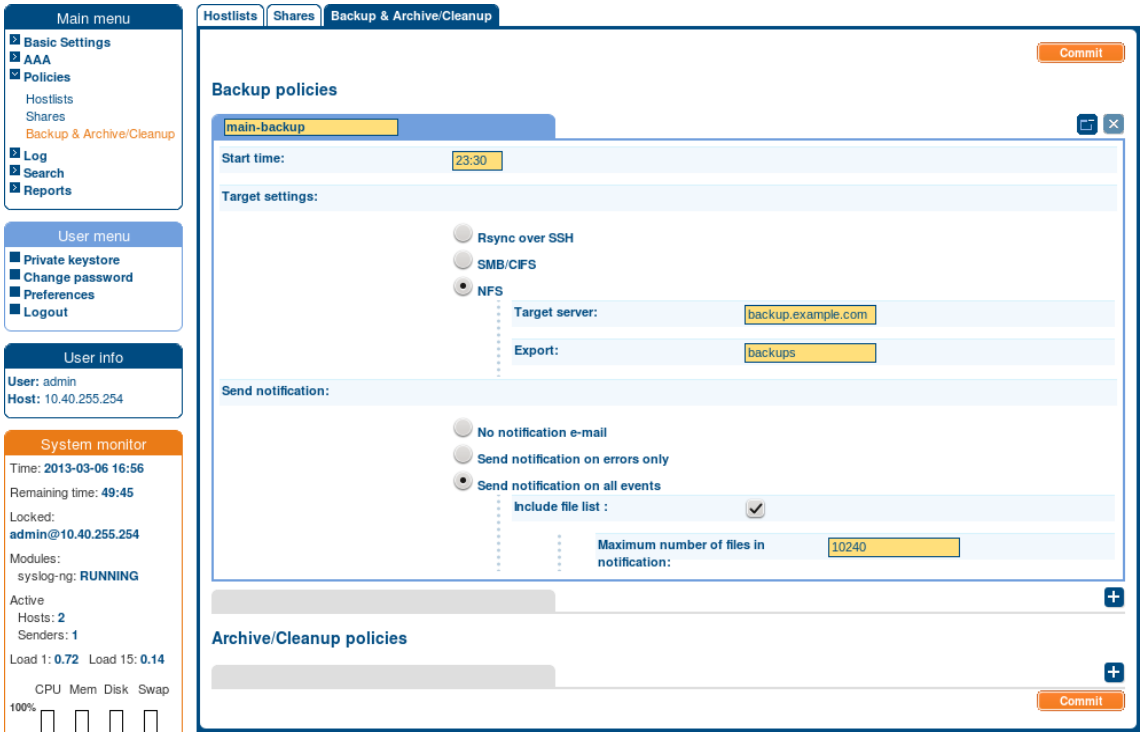

Figure 4.25. Configuring NFS backups

The backup server must also be configured to accept backups from SSB. To configure NFS on the remote server, complete the following steps:

Note
These steps must be performed on the remote server, not on SSB.

Steps:

Step 1. Add a line that corresponds to the settings of SSB to the `/etc/exports` file of the backup server. This line should contain the following parameters:

- The path to the backup directory as set in the **Export** field of the SSB backup or archiving policy.
- The IP address of the SSB interface that is used to access the remote server (that is, the address of the external interface, or the address of the management interface if it is enabled and the routing table of SSB is correctly configured — for details, see *Section 4.3, Network settings (p. 36)*.
- The following parameters: *`(rw,no_root_squash,sync)`*.

Example 4.4. Configuring NFS on the remote server
For example, if SSB connects the remote server from the *`192.168.1.15`* IP address and the data is saved into the *`/var/backups/SSB`* directory, add the following line to the `/etc/exports` file:

```
/var/backups/SSB 192.168.1.15(rw,no_root_squash,sync)
```

Step 2. Execute the following command: `exportfs -a`

Step 3. Verify that the *`rpc portmapper`* and *`rpc.statd`* applications are running.

4.9.4. Ownership of the backup files

The different backup protocols assign different file ownerships to the files saved on the backup server. The owners of the backup files created using the different protocols are the following:

- *rsync*: The user provided on the web interface.
- *SMB*: The user provided on the web interface.
- *NFS*: *`root`* with *`no-root-squash`*, *`nobody`* otherwise.

Warning
SSB cannot modify the ownership of a file that already exists on the remote server. If you change the backup protocol but you use the same directory of the remote server to store the backups, make sure to adjust the ownership of the existing files according to the new protocol. Otherwise SSB cannot overwrite the files and the backup procedure fails.

Chapter 5. User management and access control

The **AAA** menu (Authentication, Authorization, and Accounting) allows you to control the authentication, authorization, and accounting settings of the users accessing SSB. The following will be discussed in the next sections:

- For details on how to authenticate locally on SSB — see *Procedure 5.1, Managing SSB users locally (p. 68)*.
- For details on how to authenticate users using an external LDAP (for example Microsoft Active Directory) database — see *Procedure 5.4, Managing SSB users from an LDAP database (p. 72)*.
- For details on how to authenticate users using an external RADIUS server — see *Procedure 5.5, Authenticating users to a RADIUS server (p. 76)*.
- For details on how to control the privileges of users and usergroups — see *Section 5.6, Managing user rights and usergroups (p. 77)*.
- For details on how to display the history of changes of SSB configuration — see *Section 5.7, Listing and searching configuration changes (p. 82)*.

5.1. Procedure – Managing SSB users locally

Purpose:

By default, SSB users are managed locally on SSB. To create and delete local users, modify the group membership of local users, or to modify the password of a user, complete the following procedure.

Note
The `admin` user is available by default and has all possible privileges. It is not possible to delete this user.

Local users cannot be managed when LDAP authentication is used (see *Procedure 5.4, Managing SSB users from an LDAP database (p. 72)*). When LDAP authentication is enabled, the accounts of local users is disabled, they are not displayed on the **AAA** > **Local Users** page, but they are not deleted,

When using RADIUS authentication together with local users, the users are authenticated to the RADIUS server, only their group memberships must be managed locally on SSB. For details, see *Procedure 5.5, Authenticating users to a RADIUS server (p. 76)*.

Steps:

Step 1. Navigate to **AAA** > **Local Users** and click ⊞.

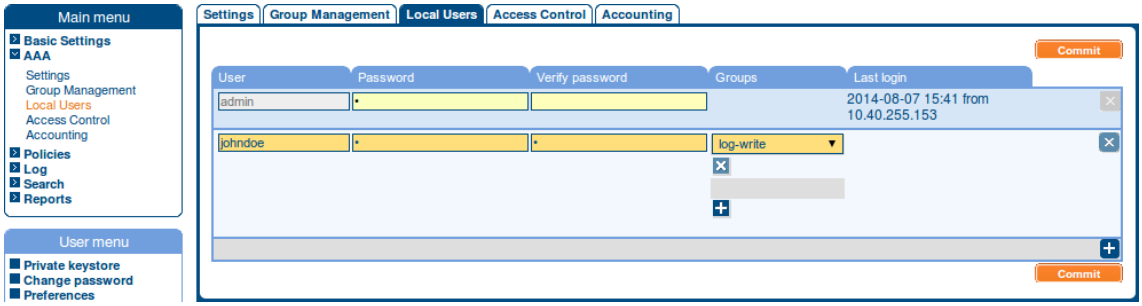

Figure 5.1. Creating local users

Step 2. Enter the username into the **User** field.

Note
The following characters cannot be used in usernames: `\/[]:;|=,+*?<>`

Step 3. Enter a password for the user into the **Password** and **Verify password** fields.

Note
SSB accepts passwords that are not longer than 150 characters. The following special characters can be used: `!"#$%&'()*+,-./:;<=>?@[\]^-`{|}`

The strength of the password is indicated below the **Password** field as you type. To set a policy for password strength, see *Procedure 5.2, Setting password policies for local users (p. 69)*. The user can change the password later from the SSB web interface.

Step 4. Click ⊞ in the **Groups** section and select a group that the user will be member of. Repeat this step to add the user to multiple groups. For details about the different groups, see *Section 5.6, Managing user rights and usergroups (p. 77)*.

- To remove a user from a group, click ⊠ next to the group.
- To delete a user, click ⊠ at the right edge of the screen.

Step 5. Click Commit.

5.2. Procedure – Setting password policies for local users

Purpose:

SSB can use password policies to enforce minimal password strength and password expiry. To create a password policy, complete the following steps.

Note
Password policies apply only for locally managed users, it has no effect if you manage your users from an LDAP database, or if you authenticate your users to a RADIUS server.

Password policies do not apply to the built-in *admin* user.

Steps:

Step 1. Navigate to **AAA > Settings**.

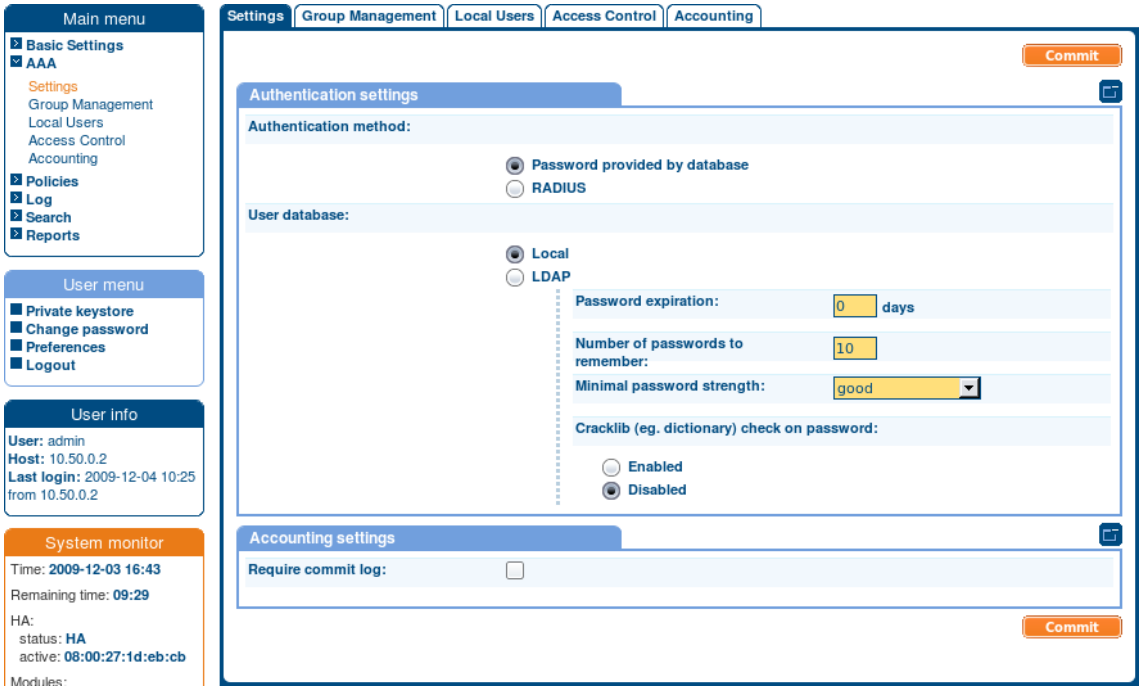

Figure 5.2. Configuring password policies

Step 2. Verify that the **Authentication method** is set to **Password provided by database** and that the **User database** is set to **Local**.

Note
If the setting of these fields is different (for example LDAP or RADIUS), then SSB is not configured to manage passwords locally.

Step 3. Set how long the passwords are valid in the **Password expiration** field. After this period, SSB users will have to change their password. To disable password expiry, enter *0*.

Step 4. To prevent password-reuse (for example when a user has two password and instead of changing to a new password only switches between the two), set how many different passwords must the user use before reusing an old password.

Step 5. To enforce the use of strong password, select the level of password-complexity from the **Minimal password strength** field.

Note
The strength of the password is determined by its entropy: the variety of numbers, letters, capital letters, and special characters used, not only by its length.

The **Enable cracklib** option executes some simple dictionary-based attacks to find weak passwords.

Step 6. Click Commit.

Note
Changes to the password policy do not affect existing passwords. However, setting password expiry will require every user to change their passwords after the expiry date, and the new passwords must comply with the strength requirements set in the password policy.

5.3. Procedure – Managing local usergroups

Purpose:

To display which users belong to a particular local usergroup, navigate to **AAA > Group Management**. You can edit the group memberships here as well.

You can use local groups to control the privileges of SSB local and LDAP users — who can view and configure what. Local groups can be also used to control access to the logfiles available via a shared folder. For details, see *Section 8.6, Accessing log files across the network (p. 150)*.

For the description of built-in groups, see *Section 5.6.5, Built-in usergroups of SSB (p. 81)*. To create a new group, complete the following steps:

Steps:

Step 1. Navigate to **AAA > Group Management** and click +.

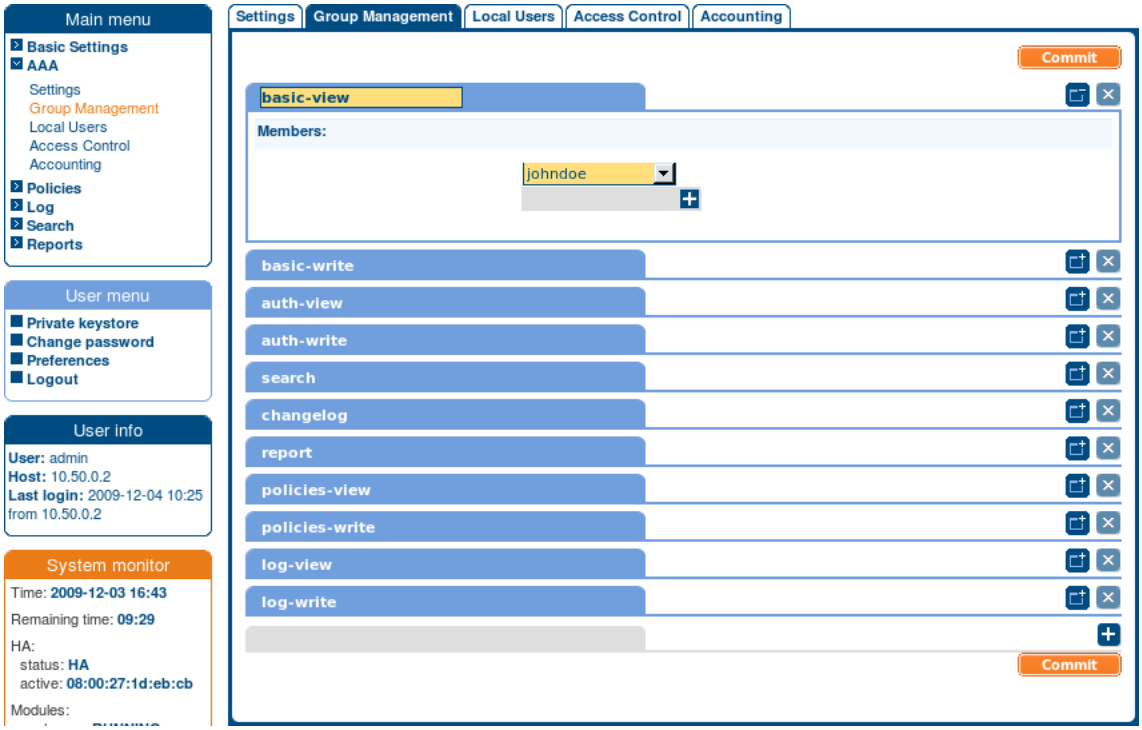

Figure 5.3. Group management

Step 2. Enter a name for the group.

Step 3. Enter the names of the users belonging to the group. Click ⊞ to add more users.

Step 4. Click Commit.

5.4. Procedure – Managing SSB users from an LDAP database

Purpose:

The SSB web interface can authenticate users to an external LDAP database to simplify the integration of SSB to your existing infrastructure. You can also specify multiple LDAP servers; if the first server is unavailable, SSB will try to connect to the second server. To enable LDAP authentication, complete the following steps.

Note

The `admin` user is available by default and has all privileges. It is not possible to delete this user.

The `admin` user can login to SSB even if LDAP authentication is used.

Enabling LDAP authentication automatically disables the access of every local user except for `admin`.

SSB accepts both pre-win2000-style and Win2003-style account names (User Principal Names). User Principal Names (UPNs) consist of a username, the at (@) character, and a domain name, for example `administrator@example.com`.

The following characters cannot be used in usernames and group names: `/\[];|=,+*)?<>@"`

When using RADIUS authentication together with LDAP users, the users are authenticated to the RADIUS server, only their group memberships must be managed in LDAP. For details, see *Procedure 5.5, Authenticating users to a RADIUS server (p. 76)*.

Warning
A user can belong to a maximum of 10,000 groups; further groups are ignored.

Warning
By default, SSB uses nested groups when querying the LDAP server. Nested groups are mostly useful when authenticating the users to Microsoft Active Directory, but can slow down the query and cause the connection to timeout if the LDAP tree is very large. In this case, disable the **Enable nested groups** option.

Steps:

Step 1. Navigate to **AAA > Settings > Authentication settings**.

Step 2. Select the **LDAP** option and enter the parameters of your LDAP server.

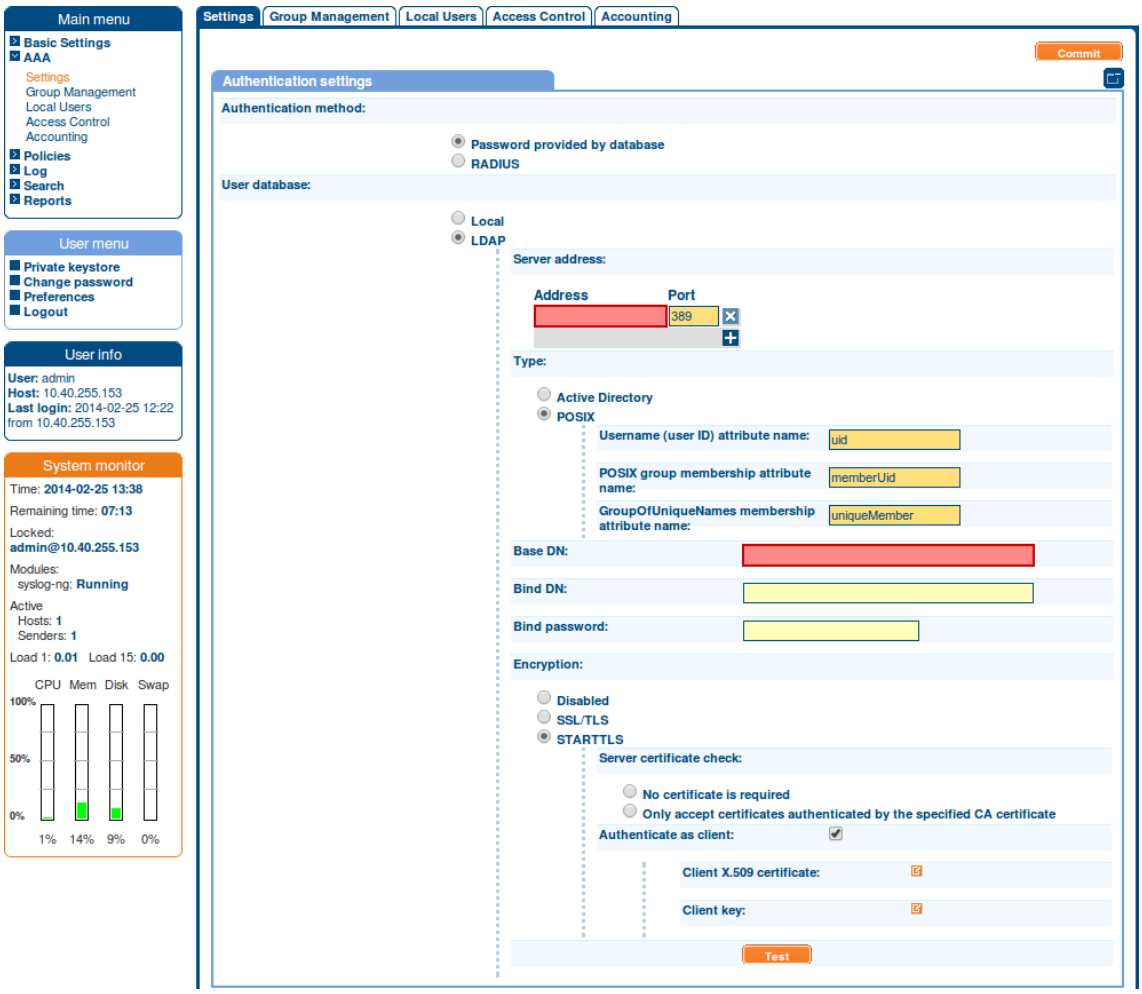

Figure 5.4. Configuring LDAP authentication

Step a. Enter the IP address or hostname and port of the LDAP server into the **Server Address** field. If you want to encrypt the communication between SSB and the LDAP server,

in case of SSL/TLS, enter 636 as the port number, or in case of STARTTLS, enter 389 as the port number.

To add multiple servers, click ⊕ and enter the address of the next server. If a server is unreachable, SSB will try to connect to the next server in the list in failover fashion.

Warning
If you will use a TLS-encrypted with certificate verification to connect to the LDAP server, use the full domain name (for example *`ldap.example.com`*) in the **Server Address** field, otherwise the certificate verification might fail. The name of the LDAP server must appear in the Common Name of the certificate.

Step b. Select the type of your LDAP server in the **Type** field. Select **Active Directory** to connect to Microsoft Active Directory servers, or **Posix** to connect to servers that use the POSIX LDAP scheme.

Step c. Enter the name of the DN to be used as the base of the queries into the **Base DN** field (for example *`DC=demodomain,DC=exampleinc`*).

Step d. Enter the name of the DN where SSB should bind to before accessing the database into the **Bind DN** field.

For example: *`CN=Administrator,CN=Users,DC=demodomain,DC=exampleinc`*.

Note
SSB accepts both pre-win2000-style and Win2003-style account names (User Principal Names), for example *`administrator@example.com`* is also accepted.

Note
Do not use sAMAccountName, as the bind DN expects a CN.

Step e. Enter the password to use when binding to the LDAP server into the **Bind Password** field.

Note
SSB accepts passwords that are not longer than 150 characters. The following special characters can be used: *`!"#$%&'()*+,-./:;<=>?@[\]^-`{|}`*

Step 3. If you want to encrypt the communication between SSB and the LDAP server, in **Encryption**, select the **SSL/TLS** or the **STARTTLS** option and complete the following steps:

Note
TLS-encrypted connection to Microsoft Active Directory is supported only on Windows 2003 Server and newer platforms. Windows 2000 Server is not supported.

- If you want SSB to verify the certificate of the server, select **Only accept certificates authenticated by the specified CA certificate** and click the icon in the **CA X.509 certificate** field. A popup window is displayed.

  Click **Browse**, select the certificate of the Certificate Authority (CA) that issued the certificate of the LDAP server, then click **Upload**. Alternatively, you can paste the certificate into the **Copy-paste** field and click **Set**.

  SSB will use this CA certificate to verify the certificate of the server, and reject the connections if the verification fails.

Warning
If you will use a TLS-encrypted with certificate verification to connect to the LDAP server, use the full domain name (for example *`ldap.example.com`*) in the **Server Address** field, otherwise the certificate verification might fail. The name of the LDAP server must appear in the Common Name of the certificate.

- If the LDAP server requires mutual authentication, that is, it expects a certificate from SSB, enable **Authenticate as client**. Generate and sign a certificate for SSB, then click in the **Client X.509 certificate** field to upload the certificate. After that, click in the **Client key** field and upload the private key corresponding to the certificate.

SSB accepts private keys in PEM (RSA and DSA), PUTTY, and SSHCOM/Tectia format. Password-protected private keys are also supported.

BalaBit recommends using 2048-bit RSA keys (or stronger).

Step 4. *Optional Step*: If your LDAP server uses a custom POSIX LDAP scheme, you might need to set which LDAP attributes store the username, or the attributes that set group memberships. For example, if your LDAP scheme does not use the *`uid`* attribute to store the usernames, set the **Username (userid) attribute name** option. You can customize group-membership attributes using the **POSIX group membership attribute name** and **GroupOfUniqueNames membership attribute name** options.

Step 5. Click Commit.

Note
You also have to configure the usergroups in SSB and possibly in your LDAP database. For details on using usergroups, see *Section 5.6.4, How to use usergroups (p. 80)*.

Step 6. Click **Test** to test the connection. Note that the testing of SSL-encrypted connections is currently not supported.

5.5. Procedure – Authenticating users to a RADIUS server

Purpose:

SSB can authenticate its users to an external RADIUS server. Group memberships of the users must be managed either locally on SSB or in an LDAP database.

Warning
The challenge/response authentication methods is currently not supported. Other authentication methods (for example password, SecureID) should work.

To authenticate SSB users to a RADIUS server, complete the following steps:

Steps:

Step 1. Navigate to **AAA > Settings**.

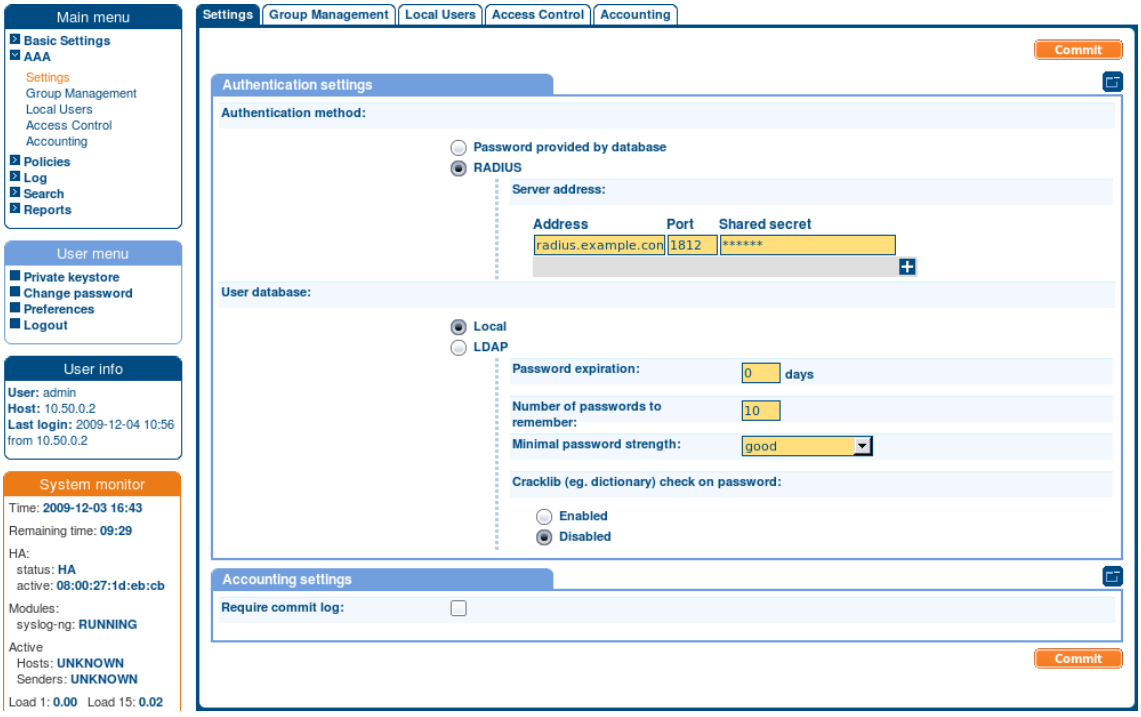

Figure 5.5. Configuring RADIUS authentication

Step 2. Set the **Authentication method** field to **RADIUS**.

Step 3. Enter the IP address or domain name of the RADIUS server into the **Address** field.

Step 4. Enter the password that SSB can use to access the server into the **Shared secret** field.

Note
SSB accepts passwords that are not longer than 150 characters. The following special characters can be used: `!"#$%&'()*+,-./:;<=>?@[\]^-`{|}`

Step 5. To add more RADIUS servers, click ⊞ and repeat Steps 2-4.
Repeat this step to add multiple servers. If a server is unreachable, SSB will try to connect to the next server in the list in failover fashion.

Step 6. When configuring RADIUS authentication with a local user database, complete the following steps.

Step a. Set **Password expiration** to `0`.

Step b. Set **Number of passwords to remember** to `0`.

Step c. Set **Minimal password strength** to `disabled`.

Step d. Set **Cracklib check on password** to `disabled`.

Step 7.

Warning
After clicking Commit, the SSB web interface will be available only after successfully authenticating to the RADIUS server. Note that the default `admin` account of SSB will be able to login normally, even if the RADIUS server is unaccessible.

Click Commit.

5.6. Managing user rights and usergroups

In SSB, user rights can be assigned to usergroups. SSB has numerous usergroups defined by default, but custom user groups can be defined as well. Every group has a set of privileges: which pages of the SSB web interface it can access, and whether it can only view (read) or also modify (read & write/perform) those pages or perform certain actions.

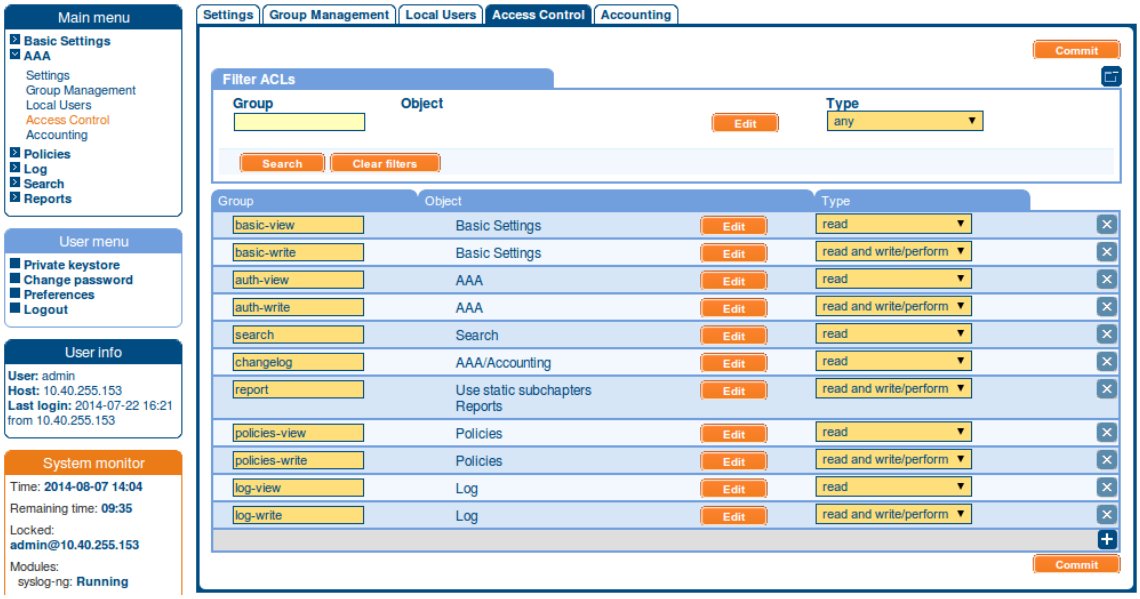

Figure 5.6. Managing SSB users

Note
Every group has either read or read & write/perform privileges to a set of pages.

- For details on modifying existing groups, see *Procedure 5.6.1, Modifying group privileges (p. 78)*.
- For details on creating a new usergroup, see *Procedure 5.6.2, Creating new usergroups for the SSB web interface (p. 79)*.
- For details on finding usergroups that have a specific privilege, see *Section 5.6.3, Finding specific usergroups (p. 80)*.
- For tips on using usergroups, see *Section 5.6.4, How to use usergroups (p. 80)*.
- For a detailed description about the privileges of the built-in usergroups, see *Section 5.6.5, Built-in usergroups of SSB (p. 81)*.

5.6.1. Procedure – Modifying group privileges

Purpose:

To modify the privileges of an existing group, complete the following steps:

Steps:

Step 1. Navigate to **AAA > Access Control**.

Step 2. Find the group you want to modify and click . The list of available privileges is displayed.

Step 3. Select the privileges (pages of the SSB interface) to which the group will have access to and click **Save**.

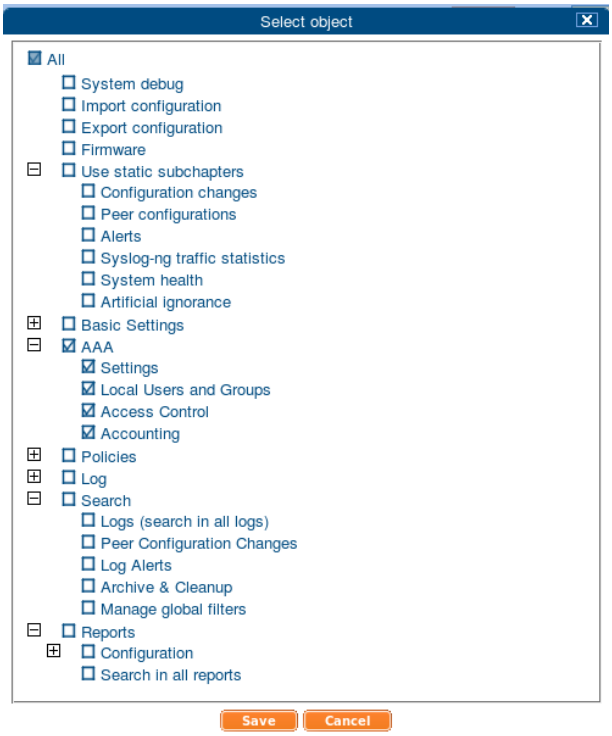

Figure 5.7. Modifying group privileges

Warning
Assigning the **Search** privilege to a user on the AAA page grants the user search access to every logspace, even if the user is not a member of the groups listed in the **Access Control** option of the particular logspace.

Step 4. Select the type of access (read or read & write) from the **Type** field.

Step 5. Click Commit.

5.6.2. Procedure – Creating new usergroups for the SSB web interface

Purpose:

To create a new group, complete the following steps:

Steps:

Step 1. Navigate to **AAA** > **Access Control** and click .

Step 2. Enter a name for the group. For details on how you should name your groups, see *Section 5.6.4, How to use usergroups (p. 80)*.

Step 3. Click located next to the name of the group. The list of available privileges is displayed.

Step 4. Select the privileges (pages of the SSB interface) to which the group will have access to and click **Save**.

Note
To export the configuration of SSB, the **Export configuration** privilege is required.

To import a configuration to SSB, the **Import configuration** privilege is required.

To update the firmware and set the active firmware, the **Firmware** privilege is required.

Step 5. Select the type of access (read or read & write) from the **Type** field.

Step 6. Click Commit.

The *admin* user is available by default and has all privileges, except that it cannot remotely access the shared logspaces. It is not possible to delete this user.

5.6.3. Finding specific usergroups

The **Filter ACLs** section of the **AAA > Access Control** page provides you with a simple searching and filtering interface to search the names and privileges of usergroups.

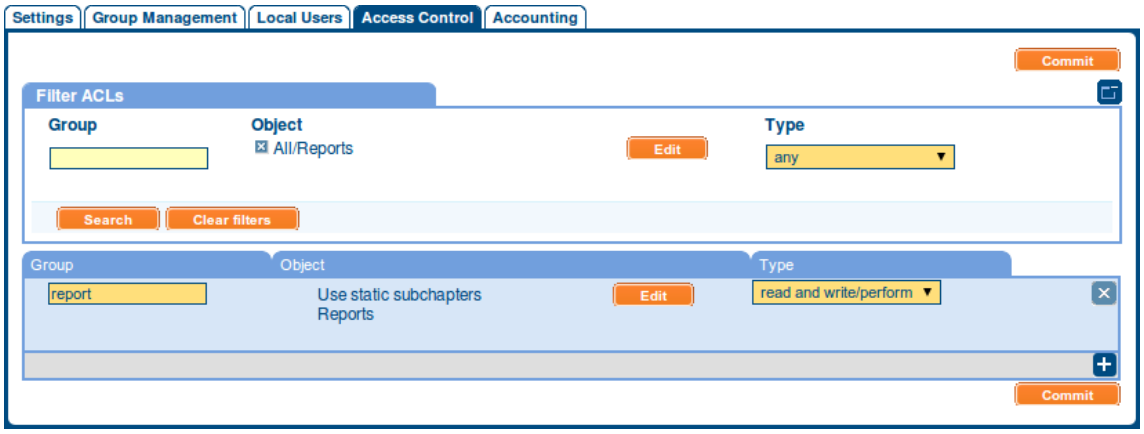

Figure 5.8. Finding specific usergroups

- To select usergroups starting with a specific string, enter the beginning of the name of the group into the **Group** field and select **Search**.
- To select usergroups who have a specific privilege, click , select the privilege or privileges you are looking for, and click **Search**.
- To filter for read or write access, use the **Type** option.

5.6.4. How to use usergroups

How you should name usergroups depends on the way you manage your SSB users.

- *Local users*: If you use only local users, create or modify your usergroups on the **AAA > Access Control** page and add users to the groups on the **AAA > Local Users** or the **AAA > Group Management** page.

- *LDAP users and LDAP groups*: If you manage your users from LDAP, and also have LDAP groups that match the way you want to group your SSB users, create or modify your usergroups on the **AAA > Access Control** page and ensure that the name of your LDAP group and the SSB usergroup is the same. For example, to make members of the `admins` LDAP group be able to use SSB, create a usergroup called `admins` on the **AAA > Access Control** page and edit the privileges of the group as needed.

Warning
A user can belong to a maximum of 10,000 groups; further groups are ignored.

- *RADIUS users and local groups*: This is the case when you manage users from RADIUS, but you cannot or do not want to create groups in LDAP. Create your local groups on the **AAA > Access Control** page, and add your RADIUS users to these groups on the **AAA > Group Management** page.

5.6.5. Built-in usergroups of SSB

SSB has the following usergroups by default:

Warning
If you use LDAP authentication on the SSB web interface and want to use the default usergroups, you have to create these groups in your LDAP database and assign users to them. For details on using usergroups, see *Section 5.6.4, How to use usergroups (p. 80)*.

- **basic-view**: View the settings in the **Basic Settings** menu, including the system logs of SSB. Members of this group can also execute commands on the **Troubleshooting** tab.

- **basic-write**: Edit the settings in the **Basic Settings** menu. Members of this group can manage SSB as a host.

- **auth-view**: View the names and privileges of the SSB administrators, the configured usergroups, and the authentication settings in the **AAA** menu. Members of this group can also view the history of configuration changes.

- **auth-write**: Edit authentication settings and manage users and usergroups.

Warning
Members of the *auth-write* group, or any other group with write privileges to the **AAA** menu are essentially equivalent to system administrators of SSB, because they can give themselves any privilege. Users with limited rights should never have such privileges.

If a user with write privileges to the **AAA** menu gives himself new privileges (for example gives himself group membership to a new group), then he has to relogin to the SSB web interface to activate the new privilege.

- **search**: Browse and download various logs and alerts in the **Search** menu.

Note
The *admin* user is not a member of this group by default, so it cannot remotely access the shared logspaces.

- **changelog**: View the history of SSB configuration changes in the **AAA > Accounting** menu.
- **report**: Browse, create and manage reports, and add statistics-based chapters to the reports in the **Reports** menu.

Note
To control exactly which statistics-based chapters and reports can the user include in a report, use the *Use static subchapters* privileges.

- **policies-view**: View the policies and settings in the **Policies** menu.
- **policies-write**: Edit the policies and settings in the **Policies** menu.

Warning
Members of this group can make the logs stored on SSB available as a shared network drive. In case of unencrypted logfiles, this may result in access to sensitive data.

- **log-view**: View the logging settings in the **Log** menu.
- **log-write**: Configure logging settings in the **Log** menu.

5.7. Listing and searching configuration changes

SSB automatically tracks every change of its configuration. To display the history of changes, select **AAA > Accounting**. The changes are organized as log messages, and can be browsed and searched using the regular

SSB search interface (for details, see *Chapter 12, Browsing log messages (p. 181)*). The following information is displayed about each modification:

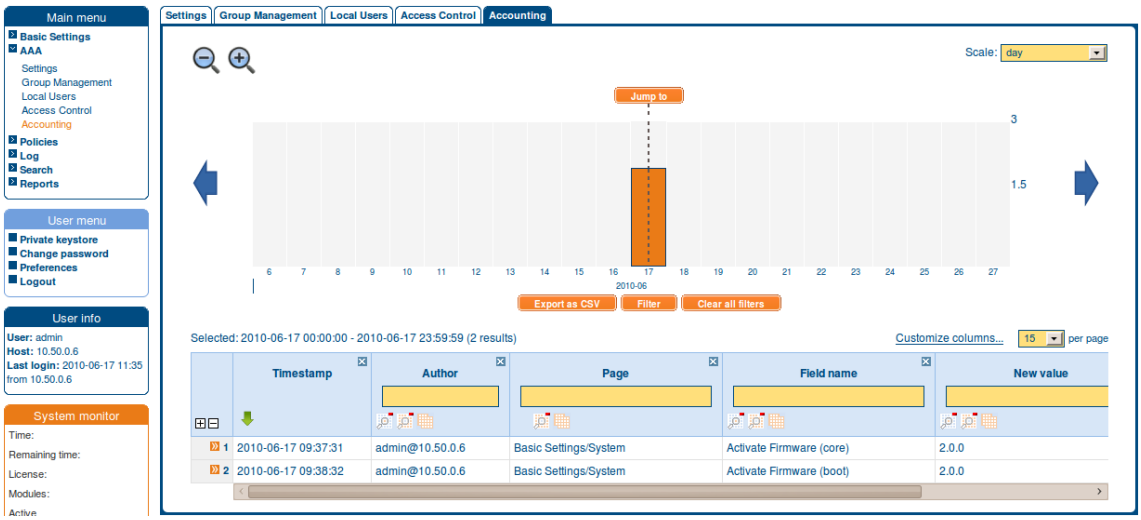

Figure 5.9. Browsing configuration changes

- **Timestamp**: The date of the modification.
- **Author**: Username of the administrator who modified the configuration of SSB.
- **Page**: The menu item that was modified.
- **Field name**: The name of the field or option that was modified.
- **New value**: The new value of the configuration parameter.
- **Message**: The changelog or commit log that the administrator submitted. This field is available only if the **Require commit log** option is enabled (see below).
- **Old value**: The old value of the configuration parameter.

To request the administrators to write an explanation to every configuration change, navigate to **AAA > Settings > Accounting settings** and select the **Require commit log** option.

Chapter 6. Managing SSB

The following sections explain the basic management tasks of SSB.

- For basic management tasks (reboot and shutdown, disabling traffic), see *Section 6.1, Controlling SSB — restart, shutdown (p. 84)*.
- For managing a high availability cluster, see *Section 6.2, Managing a high availability SSB cluster (p. 85)*.
- For instructions on upgrading SSB, see *Section 6.3, Upgrading SSB (p. 94)*.
- For instructions on accessing SSB through console and SSH, see *Section 6.4, Accessing the SSB console (p. 103)*.
- For enabling sealed mode (which disables basic configuration changes from a remote host), see *Section 6.5, Sealed mode (p. 107)*.
- For information on configuring the out-of-band (IPMI) interface, see *Section 6.6, Out-of-band management of SSB (p. 107)*.
- For managing certificates used on SSB, see *Section 6.7, Managing the certificates used on SSB (p. 110)*.
- For creating hostlist policies, see *Section 6.8, Creating hostlist policies (p. 122)*.

6.1. Controlling SSB — restart, shutdown

To restart or shut down SSB, navigate to **Basic Settings > System > System control > This node** and click the respective action button. The **Other node** refers to the slave node of a high availability SSB cluster. For details on high availability clusters, see *Section 6.2, Managing a high availability SSB cluster (p. 85)*.

Warning

- When rebooting the nodes of a cluster, reboot the other (slave) node first to avoid unnecessary takeovers.
- When shutting down the nodes of a cluster, shut down the other (slave) node first. When powering on the nodes, start the master node first to avoid unnecessary takeovers.
- When both nodes are running, avoid interrupting the connection between the nodes: do not unplug the Ethernet cables, reboot the switch or router between the nodes (if any), or disable the HA interface of SSB.

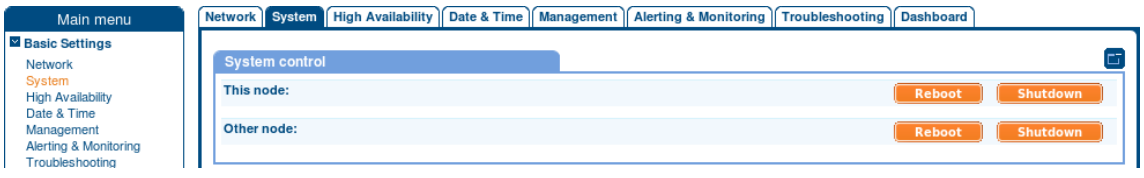

Figure 6.1. Performing basic management

Note

Web sessions to the SSB interface are persistent and remain open after rebooting SSB, so you do not have to relogin after a reboot.

6.2. Managing a high availability SSB cluster

High availability (HA) clusters can stretch across long distances, such as nodes across buildings, cities or even continents. The goal of HA clusters is to support enterprise business continuity by providing location-independent failover and recovery.

To set up a high availability cluster, connect two SSB units with identical configurations in high availability mode. This creates a master-slave (active-backup) node pair. Should the master node stop functioning, the slave node takes over the MAC addresses of the master node's interfaces. This way, the SSB servers are continuously accessible.

Note
To use the management interface and high availability mode together, connect the management interface of both SSB nodes to the network, otherwise you will not be able to access SSB remotely when a takeover occurs.

The master node shares all data with the slave node using the HA network interface (labeled as *4* or *HA* on the SSB appliance). The disks of the master and the slave node must be synchronized for the HA support to operate correctly. Interrupting the connection between running nodes (unplugging the Ethernet cables, rebooting a switch or a router between the nodes, or disabling the HA interface) disables data synchronization and forces the slave to become active. This might result in data loss. You can find instructions to resolve such problems and recover an SSB cluster in *Section 16.6, Troubleshooting an SSB cluster (p. 228)*.

The **Basic Settings > High Availability** page provides information about the status of the HA cluster and its nodes.

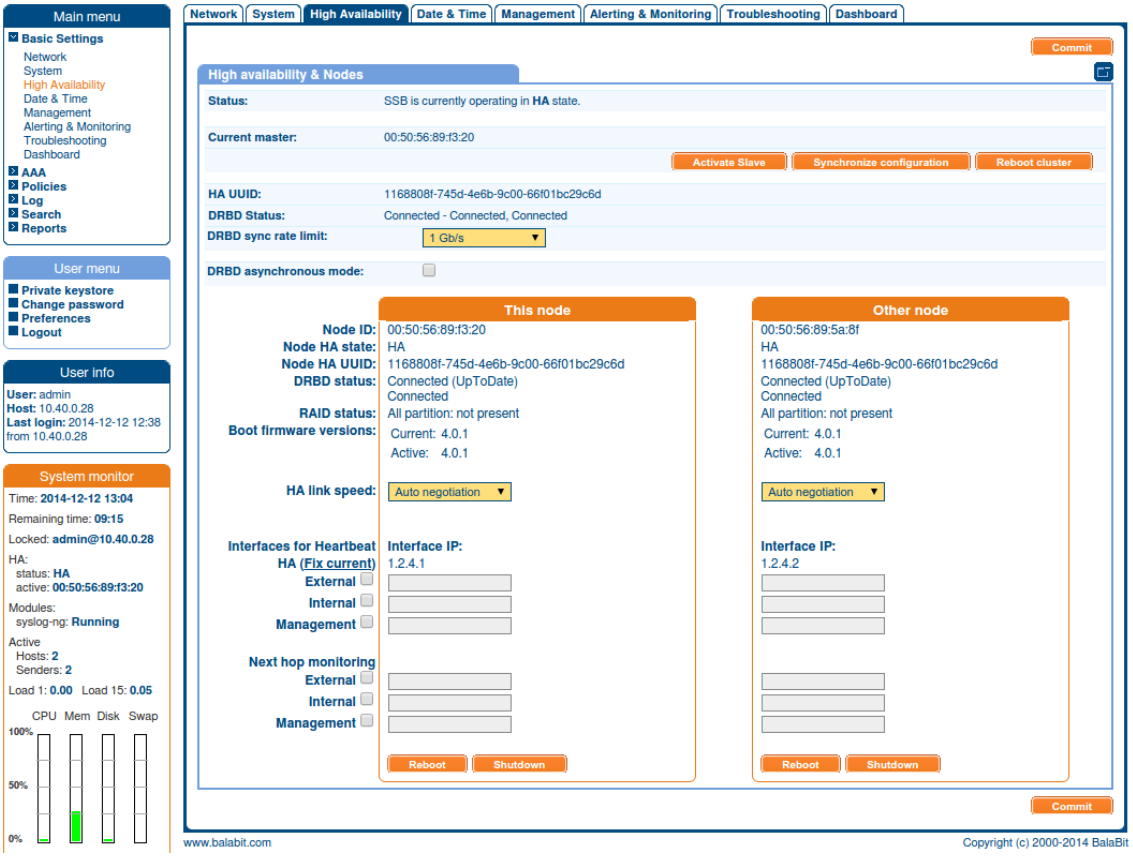

Figure 6.2. Managing a high availability cluster

The following information is available about the cluster:

- **Status**: Indicates whether the SSB nodes recognize each other properly and whether those are configured to operate in high availability mode.

  You can find the description of each HA status in *Section 16.6.1, Understanding SSB cluster statuses (p. 228)*.

- **Current master**: The MAC address of the high availability interface (*4* or *HA*) of the node. This address is also printed on a label on the top cover of the SSB unit.

- **HA UUID**: A unique identifier of the HA cluster. Only available in High Availability mode.

- **DRBD status**: Indicates whether the SSB nodes recognize each other properly and whether those are configured to operate in high availability mode.

  You can find the description of each DRBD status in *Section 16.6.1, Understanding SSB cluster statuses (p. 228)*.

- **DRBD sync rate limit**: The maximum allowed synchronization speed between the master and the slave node.

You can find more information about configuring the DRBD sync rate limit in *Section 6.2.1, Adjusting the synchronization speed (p. 89)*.

The active (master) SSB node is labeled as **This node**, this unit receives the incoming log messages and provides the web interface. The SSB unit labeled as **Other node** is the slave node that is activated if the master node becomes unavailable.

The following information is available about each node:

- **Node ID**: The MAC address of the *HA* interface of the node. This address is also printed on a label on the top cover of the SSB unit.

  For SSB clusters, the IDs of both nodes are included in the internal log messages of SSB.

- **Node HA state**: Indicates whether the SSB nodes recognize each other properly and whether those are configured to operate in high availability mode.

  You can find the description of each HA status in *Section 16.6.1, Understanding SSB cluster statuses (p. 228)*.

- **Node HA UUID**: A unique identifier of the cluster. Only available in High Availability mode.

- **DRBD status**: The status of data synchronization between the nodes.

  You can find the description of each DRBD status in *Section 16.6.1, Understanding SSB cluster statuses (p. 228)*.

- **Raid status**: The status of the RAID device of the node.

- **Boot firmware version**: Version number of the boot firmware.

  You can find more information about the boot firmware in *Section 2.8, Firmware in SSB (p. 9)*.

- **HA link speed**: The maximum allowed speed between the master and the slave node. The HA link's speed must exceed the **DRBD sync rate limit**, else the web UI might become unresponsive and data loss can occur.

- **Interfaces for Heartbeat**: Virtual interface used only to detect that the other node is still available; it is not used to synchronize data between the nodes (only heartbeat messages are transferred).

  You can find more information about configuring redundant heartbeat interfaces in *Procedure 6.2.3, Redundant heartbeat interfaces (p. 90)*.

- **Next hop monitoring**: IP addresses (usually next hop routers) to continuously monitor from both the master and the slave nodes using ICMP echo (ping) messages. If any of the monitored addresses becomes unreachable from the master node while being reachable from the slave node (in other words, more monitored addresses are accessible from the slave node) then it is assumed that the master node is unreachable and a forced takeover occurs – even if the master node is otherwise functional.

You can find more information about configuring next-hop monitoring in *Procedure 6.2.4, Next-hop router monitoring (p. 92)*.

The following configuration and management options are available for HA clusters:

- *Set up a high availability cluster:* You can find detailed instructions for setting up a HA cluster in *Procedure B.2, Installing two SSB units in HA mode (p. 240)*.
- *Adjust the DRBD (master-slave) synchronization speed:* You can change the limit of the DRBD synchronization rate.
  You can find more information about configuring the DRBD synchronization speed in *Section 6.2.1, Adjusting the synchronization speed (p. 89)*.

- *Enable asynchronous data replication:* You can compensate for high network latency and bursts of high activity by enabling asynchronous data replication between the master and the slave node with the **DRBD asynchronous mode** option.

  You can find more information about configuring asynchronous data replication in *Section 6.2.2, Asynchronous data replication (p. 89)*.

- *Configure redundant heartbeat interfaces:* You can configure virtual interfaces for each HA node to monitor the availability of the other node.
  You can find more information about configuring redundant heartbeat interfaces in *Procedure 6.2.3, Redundant heartbeat interfaces (p. 90)*.

- *Configure next-hop monitoring:* You can provide IP addresses (usually next hop routers) to continuously monitor from both the master and the slave nodes using ICMP echo (ping) messages. If any of the monitored addresses becomes unreachable from the master node while being reachable from the slave node (in other words, more monitored addresses are accessible from the slave node) then it is assumed that the master node is unreachable and a forced takeover occurs – even if the master node is otherwise functional.
  You can find more information about configuring next-hop monitoring in *Procedure 6.2.4, Next-hop router monitoring (p. 92)*.

- *Reboot the HA cluster:* To reboot both nodes, click **Reboot Cluster**. To prevent takeover, a token is placed on the slave node. While this token persists, the slave node halts its boot process to make sure that the master node boots first. Following reboot, the master removes this token from the slave node, allowing it to continue with the boot process.

  If the token still persists on the slave node following reboot, the **Unblock Slave Node** button is displayed. Clicking the button removes the token, and reboots the slave node.

- *Reboot a node:* Reboots the selected node.
  When rebooting the nodes of a cluster, reboot the other (slave) node first to avoid unnecessary takeovers.

- *Shutdown a node:* Forces the selected node to shutdown.

When shutting down the nodes of a cluster, shut down the other (slave) node first. When powering on the nodes, start the master node first to avoid unnecessary takeovers.

- *Manual takeover:* To activate the other node and disable the currently active node, click **Activate slave**.

  Activating the slave node terminates all connections of SSB and might result in data loss. The slave node becomes active after about 60 seconds, during which SSB cannot accept incoming messages. Enable disk-buffering on your syslog-ng clients and relays to prevent data loss in such cases.

6.2.1. Adjusting the synchronization speed

When operating two SSB units in High Availability mode, every incoming data copied from the master (active) node to the slave (passive) node. Since synchronizing data can take up significant system-resources, the maximal speed of the synchronization is limited, by default, to `10 MB/sec`. However, this means that synchronizing large amount of data can take very long time, so it is useful to increase the synchronization speed in certain situations — for example, when synchronizing the disks after converting a single node to a high availability cluster.

The **Basic Settings > High Availability > DRBD status** field indicates whether the latest data (including SSB configuration, log files, and so on) is available on both SSB nodes. For a description of each possible status, see *Section 16.6.1, Understanding SSB cluster statuses (p. 228)*.

To change the limit of the DRBD synchronization rate, navigate to **Basic Settings > High Availability**, select **DRBD sync rate limit**, and select the desired value.

Set the sync rate carefully. A high value is not recommended if the load of SSB is very high, as increasing the resources used by the synchronization process may degrade the general performance of SSB. On the other hand, the HA link's speed must exceed the speed of the incoming logs, else the web UI might become unresponsive and data loss can occur.

If you experience bursts of high activity, consider turning on asynchronous data replication.

6.2.2. Asynchronous data replication

When a high availability SSB cluster is operating in a high-latency environment or during brief periods of high load, there is a risk of slowness, latency or package loss. To manage this, you can compensate latency with asynchronous data replication.

Asynchronous data replication is a method where local write operations on the primary node are considered complete when the local disk write is finished and the replication packet is placed in the local TCP send buffer. It does not impact application performance, and tolerates network latency, allowing the use of physically distant storage nodes. However, because data is replicated at some point after local acknowledgement, the remote storage nodes are slightly out of step; if the local node at the primary data center breaks down, data loss occurs.

To turn asynchronous data replication on, navigate to **Basic Settings > High Availability**, and enable **DRBD asynchronous mode**. You have to reboot the cluster (click **Reboot cluster**) for the change to take effect.

Under prolonged heavy load, asynchronous data replication might not be able to compensate for latency or for high packet loss ratio (over 1%). In this situation, stopping the slave machine is recommended to avoid data loss at the temporary expense of redundancy.

6.2.3. Procedure – Redundant heartbeat interfaces

Purpose:

To avoid unnecessary takeovers and to minimize the chance of split-brain situations, you can configure additional heartbeat interfaces in SSB. These interfaces are used only to detect that the other node is still available; they are not used to synchronize data between the nodes (only heartbeat messages are transferred). For example, if the main HA interface breaks down, or is accidentally unplugged and the nodes can still access each other on the redundant HA interface, no takeover occurs, but no data is synchronized to the slave node until the main HA link is restored. Similarly, if connection on the redundant heartbeat interface is lost, but the main HA connection is available, no takeover occurs.

If a redundant heartbeat interface is configured, its status is displayed in the **Basic Settings > High Availability > Redundant Heartbeat status** field, and also in the **HA > Redundant** field of the System monitor. For a description of each possible status, see *Section 16.6.1, Understanding SSB cluster statuses (p. 228)*.

The redundant heartbeat interface is a virtual interface with a virtual MAC address that uses an existing interface of SSB (for example, the external or the management interface). The original MAC address of the interface is displayed as **Production MAC**, and the MAC address of the virtual redundant heartbeat interface is displayed as **HA MAC**.

The MAC address of the redundant heartbeat interface is generated in a way that it cannot interfere with the MAC addresses of physical interfaces. Similarly, the HA traffic on the redundant heartbeat interface cannot interfere with any other traffic on the interface used.

If the nodes lose connection on the main HA interface, and after a time the connection is lost on the redundant heartbeat interfaces as well, the slave node becomes active. However, as the master node was active for a time when no data synchronization was possible between the nodes, this results in a split-brain situation which must be resolved before the HA functionality can be restored. For details, see *Procedure 16.6.3, Recovering from a split brain situation (p. 231)*.

Note
Even if redundant HA links are configured, if the dedicated HA link fails, the slave node will not be visible on the High Availability page anymore.

To configure a redundant heartbeat interface, complete the following steps.

Steps:

Step 1. Navigate to **Basic Settings > High Availability > Interfaces for Heartbeat**.

Step 2. Select the interface you want to use as redundant heartbeat interface (for example `External`). Using an interface as a redundant heartbeat interface does not affect the original traffic of the interface.

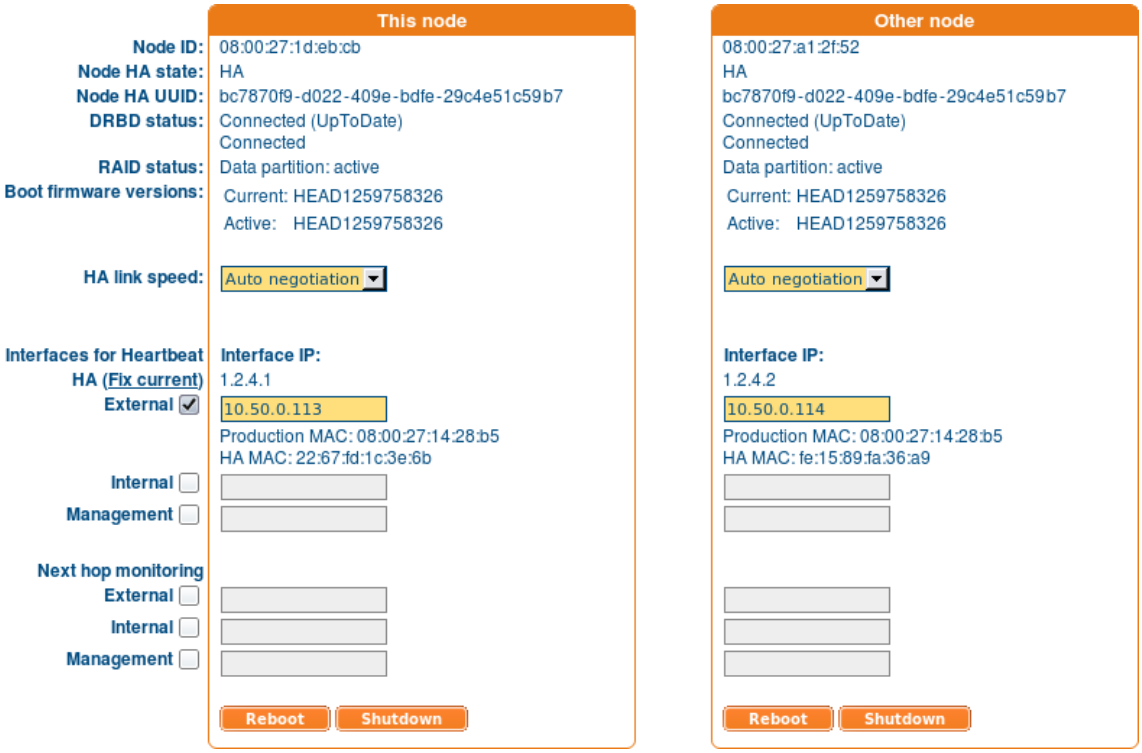

Figure 6.3. Configuring redundant heartbeat interfaces

Step 3. Enter an IP address into the **This node > Interface IP** field of the selected interface. Note the following:

- The two nodes must have different **Interface IP**.
- If you do not use next hop monitoring on the redundant interface, you can use any **Interface IP** (even if otherwise it does not exist on that network).
- If you use next hop monitoring on the redundant interface, the **Interface IP** address must be a real IP address that is visible from the other node.
- If you use next hop monitoring on the redundant interface, the **Interface IP** must be accessible from the next-hop address, and vice-versa. For details on next hop monitoring, see *Procedure 6.2.4, Next-hop router monitoring (p. 92)*.

Step 4. Enter an IP address into the **Other node > Interface IP** field of the selected interface. Note the following:

- The two nodes must have different **Interface IP**.
- If you do not use next hop monitoring on the redundant interface, you can use any **Interface IP** (even if otherwise it does not exist on that network).
- If you use next hop monitoring on the redundant interface, the **Interface IP** address must be a real IP address that is visible from the other node.

- If you use next hop monitoring on the redundant interface, the **Interface IP** must be accessible from the next-hop address, and vice-versa. For details on next hop monitoring, see *Procedure 6.2.4, Next-hop router monitoring (p. 92)*.

Step 5. Repeat the previous steps to add additional redundant heartbeat interfaces if needed.

Step 6. Click Commit.

Step 7. Restart the nodes for the changes to take effect: click **Reboot Cluster**.

6.2.4. Procedure – Next-hop router monitoring

Purpose:

By default, HA takeover occurs only if the master node stops working or becomes unreachable from the slave node. However, this does not cover the scenario when the master node becomes unaccessible to the outside world (for example its external interface or the router or switch connected to the external interface breaks down) while the slave node would be still accessible (for example because it is connected to a different router).

To address such situations, you can specify IP addresses (usually next hop routers) to continuously monitor from both the master and the slave nodes using ICMP echo (ping) messages. One such address can be set up for every interface.

When setting up next hop monitoring, you have to make sure that the master and slave nodes can ping the specified address directly. You can either:

- Choose the addresses of the redundant-HA SSB interfaces so that they are on the same subnet as the next-hop address
- Configure the next-hop device with an additional IP-address that is on the same subnet as the redundant-HA SSB interfaces facing it

If any of the monitored addresses becomes unreachable from the master node while being reachable from the slave node (in other words, more monitored addresses are accessible from the slave node) then it is assumed that the master node is unreachable and a forced takeover occurs — even if the master node is otherwise functional.

Naturally, if the slave node is not capable of taking over the master node (for example because there is data not yet synchronized from the current master node) no takeover is performed.

To configure next hop monitoring, complete the following steps.

Steps:

Step 1. Navigate to **Basic Settings > High Availability > Next hop monitoring**.

Step 2. Select the interface to use for monitoring its next-hop router.

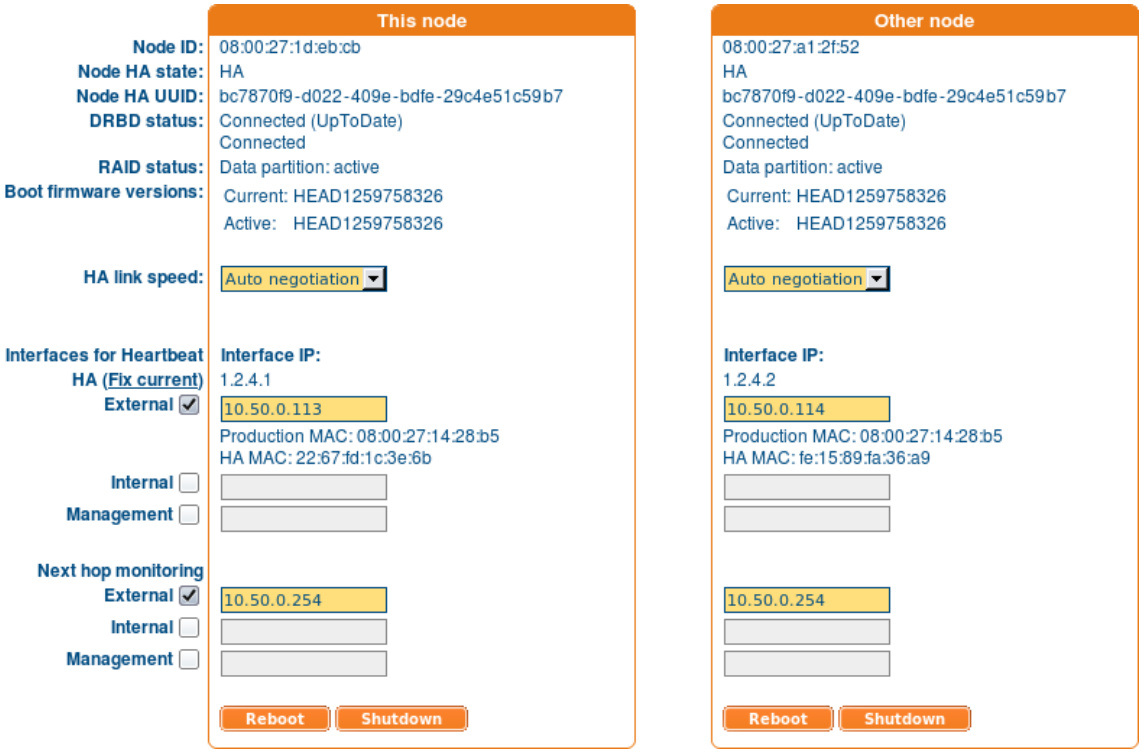

Figure 6.4. Configuring next hop monitoring

Step 3. Enter the IP address to monitor from the current master node (for example the IP address of the router or the switch connected to the interface) into the **This node > Next hop IP** field of the selected interface. This IP address must be a real IP address that is visible from the interface, and must be on the same local network segment.

Step 4. Enter the IP address to monitor from the current slave node (for example the IP address of the router or the switch connected to the interface) into the **Other node > Next hop IP** field of the selected interface. This IP address must be a real IP address that is visible from the interface, and must be on the same local network segment.

Step 5. Repeat the previous steps to add IP addresses to be monitored from the other interfaces if needed.

Step 6. Click Commit.

Warning
For the changes to take effect, you have to restart both nodes. To restart both nodes, click **Reboot Cluster**.

6.3. Upgrading SSB

SSB appliances are preinstalled with the latest available Long Term Support (LTS) release. Each LTS release is supported for 3 years after original publication date, and for 1 year after succeeding LTS Release is published (whichever date is later). You are encouraged to upgrade to succeeding LTS releases.

Feature Releases provide additional features which are not yet consolidated to an LTS release. To gain access to these features, you may install a supported Feature Release on the appliance, with the following conditions:

- You cannot roll back to an LTS release from a Feature Release.
- Feature Releases are released and supported in a timeline of 6 (+2) months. You have to keep upgrading SSB to the latest Feature Release to ensure that your appliance is supported.

For both LTS and Feature Releases, BalaBit regularly incorporates security patches and bugfixes, and issues updated Revisions of the released product. We strongly recommend always installing the latest Revision of the used software Release.

Warning
Downgrading from the latest feature release, even to an LTS release, voids support for SSB.

The following sections describe how to keep SSB up to date, and how to install a new license:

- Prerequisites: *Section 6.3.1, Upgrade checklist (p. 94)*.
- Upgrading a single node: *Procedure 6.3.2, Upgrading SSB (single node) (p. 95)*.
- Upgrading a high availability cluster: *Procedure 6.3.3, Upgrading an SSB cluster (p. 97)*.
- Troubleshooting: *Section 6.3.4, Troubleshooting (p. 98)*.
- Rollback instructions: *Procedure 6.3.5, Reverting to an older firmware version (p. 98)*.
- Renewing the SSB license: *Procedure 6.3.6, Updating the SSB license (p. 99)*.
- Exporting the configuration of SSB: *Procedure 6.3.7, Exporting the configuration of SSB (p. 100)*.
- Importing the configuration of SSB: *Procedure 6.3.8, Importing the configuration of SSB (p. 102)*.

6.3.1. Upgrade checklist

- You have created a configuration backup of SSB.
  For detailed instructions, refer to *Procedure 6.3.7, Exporting the configuration of SSB (p. 100)*.

- You have a valid MyBalaBit account.
  To download the required firmware files and the license, you need a valid MybalaBit account. You can sign up at *https://my.balabit.com/login*. Note that the registration is not automatic, and might require up to two working days to process.

- You have downloaded the latest SSB core firmware and boot firmware from *https://www.balabit.com/network-security/syslog-ng/log-server-appliance/download/syslog-ng-store-box/*.

For a detailed description of the different firmwares, see *Section 2.8, Firmware in SSB (p. 9)*.

- You have read the Release Notes of the firmware(s) before updating. The Release Notes might include additional instructions specific to the firmware version.
  The Release Notes are available here: *https://www.balabit.com/network-security/syslog-ng/log-server-appliance/support/upgrade/*.

- *If you have a high availability cluster:*
  You have IPMI access (or direct local access) to the slave node. You can find detailed information on using the IPMI interface in the following documents:

  The *Onboard BMC/IPMI User's Guide*, available on the BalaBit Hardware Documentation page at *https://www.balabit.com/support/documentation/*.

- *If you have a high availability cluster with geoclustering enabled:*
  Perform the firmware upload steps an hour before the actual upgrade. Geoclustering can introduce delays in master-slave synchronization, and the slave node might not be able to sync the new firmware from the master node on time.

During the upgrade, SSB displays information about the progress of the upgrade and any possible problems to the console, which you can monitor with IPMI (ILOM) or console access.

We recommend that you test the upgrade process in a non-productive (virtual, etc.) environment first.

Upgrading SSB requires a reboot. We strongly suggest that you perform the upgrade on the productive appliance during maintenance hours only, to avoid any potential data loss.

6.3.2. Procedure – Upgrading SSB (single node)

Steps:

Step 1. Update the core firmware of SSB using the web interface.

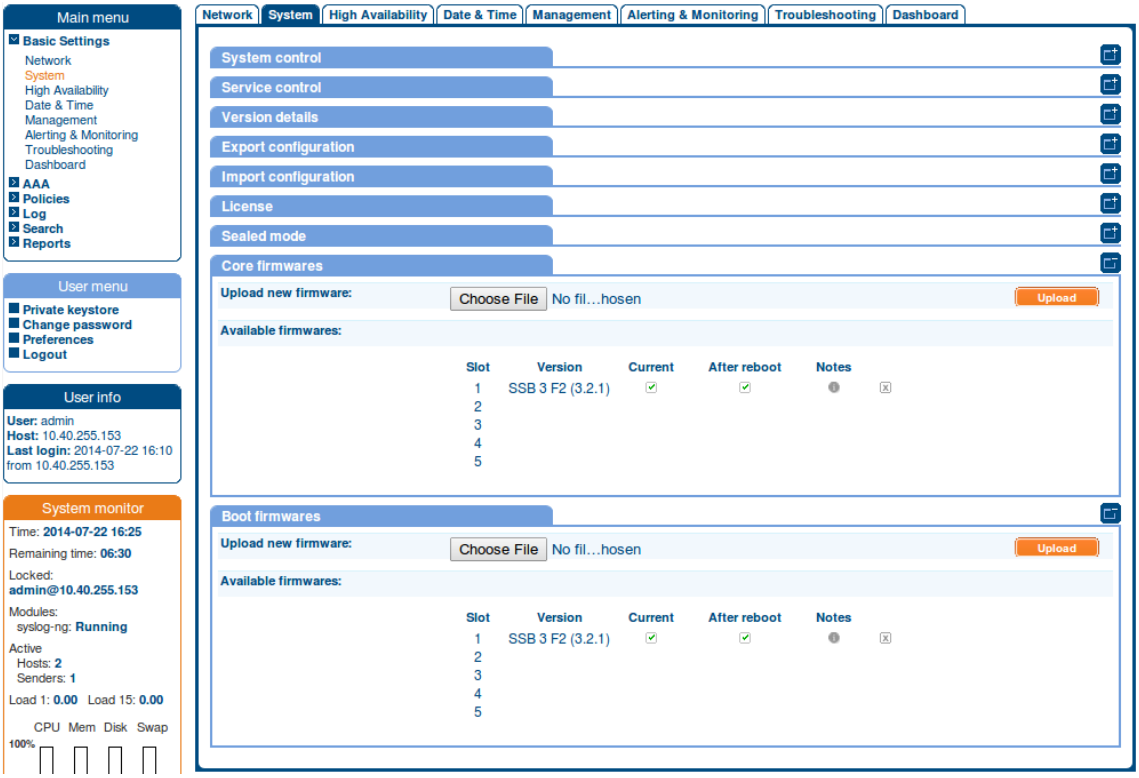

Figure 6.5. Managing the firmwares

Step a. Navigate to **Basic Settings** > **System** > **Core firmwares**.

Step b. Upload the new core firmware.

Step c. When the upload is finished, select the **After reboot** option for the new firmware.

Do not reboot SSB yet.

Step d. To read the Upgrade Notes of the uploaded firmware, click on the icon. The Upgrade Notes are displayed in a pop-up window.

Step 2. Upload the boot firmware of SSB using the web interface.

Step a. Navigate to **Basic Settings** > **System** > **Boot firmwares**.

Step b. Upload the new boot firmware.

Step c. When the upload is finished, select the **After reboot** option for the new firmware.

Step d. To read the Upgrade Notes of the uploaded firmware, click on the icon. The Upgrade Notes are displayed in a pop-up window.

Step 3. Navigate to **Basic Settings** > **System** > **System Control** > **This node**, and choose **Reboot**.

SSB attempts to boot with the new firmware. Wait for the process to complete.

Step 4. Login to the SSB web interface to verify that the upgrade was successful.

Navigate to **Basic Settings > System > Version details** and check the version numbers of SSB. In case you encounter problems, you can find common troubleshooting steps in *Section 6.3.4, Troubleshooting (p. 98)*.

6.3.3. Procedure – Upgrading an SSB cluster

Steps:

Step 1. Update the core firmware of SSB using the web interface.

Step a. Navigate to **Basic Settings > System > Core firmwares**.

Step b. Upload the new core firmware.

Step c. When the upload is finished, select the **After reboot** option for the new firmware.

Do not reboot SSB yet.

Step d. To read the Upgrade Notes of the uploaded firmware, click on the icon. The Upgrade Notes are displayed in a pop-up window.

Step 2. Upload the boot firmware of SSB using the web interface.

Step a. Navigate to **Basic Settings > System > Boot firmwares**.

Step b. Upload the new boot firmware.

Step c. When the upload is finished, select the **After reboot** option for the new firmware.

Do not reboot SSB yet.

Step d. To read the Upgrade Notes of the uploaded firmware, click on the icon. The Upgrade Notes are displayed in a pop-up window.

Step 3. Navigate to **Basic Settings > System > High availability > Other node** and click **Shutdown**.

Step 4. Restart the master node: click **This node > Reboot**.

SSB attempts to boot with the new firmware. Wait for the process to complete.

Step 5. Login to the SSB web interface to verify that the master node upgrade was successful.

Navigate to **Basic Settings > System > Version details** and check the version numbers of SSB. In case you encounter problems, you can find common troubleshooting steps in *Section 6.3.4, Troubleshooting (p. 98)*.

Step 6. Use the IPMI interface to reboot the slave node.
The slave node attempts to boot with the new firmware, and reconnect to the master node. Wait for the process to finish.

Step 7. Navigate to **Basic Settings > System > High availability** and verify that the slave node is connected, and has the same firmware versions as the master node.

6.3.4. Troubleshooting

If you experience any strange behavior of the web interface, first try to reload the page by holding the SHIFT key while clicking the Reload button of your browser to remove any cached version of the page.

In the unlikely case that SSB encounters a problem during the upgrade process and cannot revert to its original state, SSB performs the following actions:

- Initializes the network interfaces using the already configured IP addresses.
- Enables SSH-access to SSB, unless SSB is running in sealed mode. That way it is possible to access the logs of the upgrade process that helps the BalaBit Support Team to diagnose and solve the problem. Note that SSH access will be enabled on every active interface, even if management access has not been enabled for the interface.

In case the web interface is not available within 30 minutes of rebooting SSB, check the information displayed on the local console and contact the BalaBit Support Team.

6.3.5. Procedure – Reverting to an older firmware version

Purpose:

SSB can store up to five different firmware versions, any of them can be booted if required. The available firmwares are displayed on the **Basic Settings > System > Boot firmware** and **Basic Settings > System > Core firmware** pages. The list shows the detailed version of each firmware, including the version number, the revision number, and the build date. The firmware running on SSB is marked with in the **Current** column. The firmware that will be run after the next SSB reboot is marked with in the **After reboot** column.

To boot an older firmware, complete the following steps:

Warning
When upgrading SSB, it is possible that the configuration file is updated as well. In such cases, simply rebooting with the old firmware will not result in a complete downgrade, because the old firmware may not be able to read the new configuration file. If this happens, access the console menu of SSB, and select the **Revert Configuration** option to restore the configuration file to its state before the firmware was upgraded. For details on using the console menu, see *Section 6.4.1, Using the console menu of SSB (p. 103)*.

Warning
Downgrading from the latest feature release, even to an LTS release, voids support for SSB.

Steps:

Step 1. Navigate to **Basic Settings > System > Boot firmware**.

Step 2. Select the firmware version to use, and click □ in the **After reboot** column.

Step 3. Navigate to **Basic Settings > System > Core firmware**.

Step 4. Select the firmware version to use, and click □ in the **Boot** column.

Step 5. *If you are downgrading a single SSB node:*

Select **System control > This node > Reboot** to reboot SSB.

Step 6. *If you are downgrading an SSB cluster:*
Follow the instructions described in *Procedure 6.3.3, Upgrading an SSB cluster (p. 97)* (skip the upload steps). Below is a summary of the necessary actions:

Step a. You need IPMI (or direct physical) access to the slave node to proceed.

Step b. Take the slave node offline.

Step c. Restart the master node. Following reboot, verify that the master node was downgraded successfully.

Step d. Use the IPMI interface to power on the slave node. Verify that the slave node reconnected to the master, and was downgraded successfully.

6.3.6. Procedure – Updating the SSB license

Purpose:

The SSB license must be updated before the existing license expires or when you purchase a new license. Information of the current license of SSB is displayed on the **Basic Settings > System > License** page. The following information is displayed:

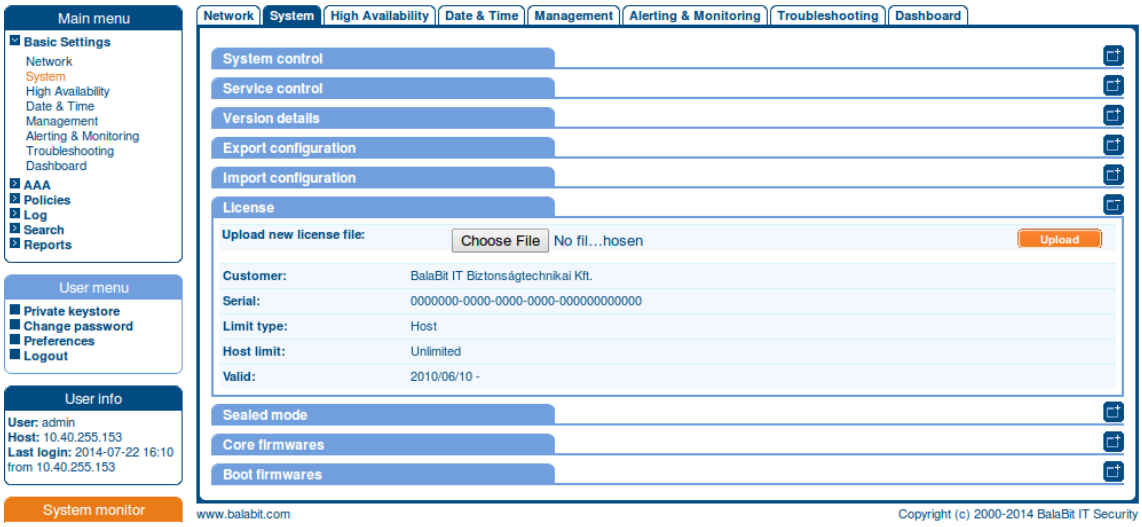

Figure 6.6. Updating the license

- **Customer**: The company permitted to use the license (for example `Example Ltd.`).

- **Serial**: The unique serial number of the license.
- **Host limit**: The number of peers SSB accepts log messages from.
- **Validity**: The period in which the license is valid. The dates are displayed in *YYYY/MM/DD* format.

SSB gives an automatic alert one week before the license expires. An alert is sent also when the number of peers exceeds 90% of the limit set in the license.

To update the license, complete the following steps:

Warning
Before uploading a new license, you are recommended to backup the configuration of SSB. For details, see *Procedure 6.3.7, Exporting the configuration of SSB (p. 100)*.

Steps:

Step 1. Navigate to **Basic Settings > System > License**.

Step 2. Click **Browse** and select the new license file.

Note
It is not required to manually decompress the license file. Compressed licenses (for example `.zip` archives) can also be uploaded.

Step 3. Click **Upload**, then Commit.

Step 4. To activate the new license, navigate to **Service control > Syslog traffic, indexing & search:** and click **Restart syslog-ng**.

6.3.7. Procedure – Exporting the configuration of SSB

Purpose:

The configuration of SSB can be exported (for manual archiving, or to migrate it to another SSB unit) from the **Basic Settings > System** page. Use the respective action buttons to perform the desired operation.

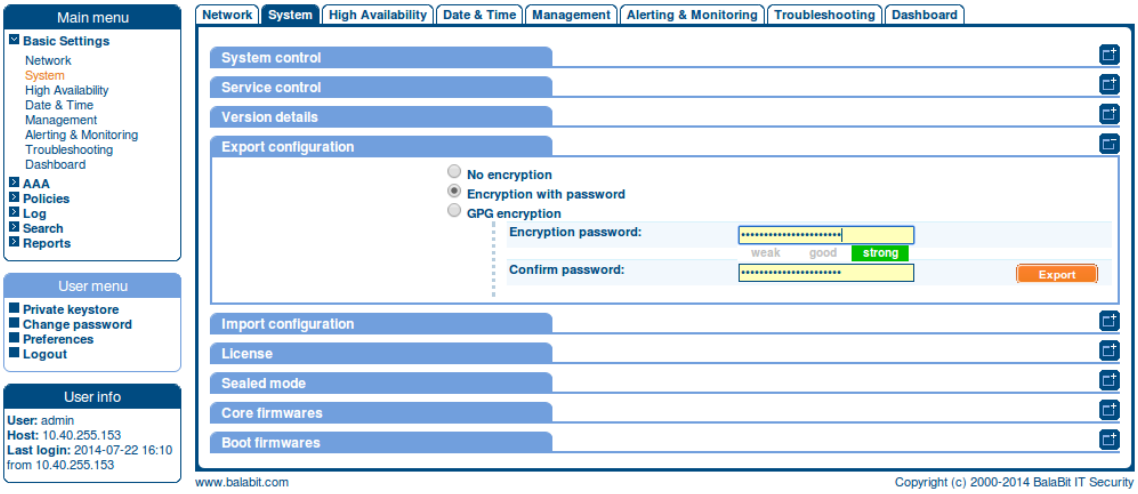

Figure 6.7. Exporting the SSB configuration

Steps:

Step 1. Navigate to **Basic Settings > System > Export configuration**.

Step 2. Select how to encrypt the configuration:

- To export the configuration file without encryption, select **No encryption**.

Warning
Exporting the SSB configuration without encyption is not recommended, as it contains sensitive information such as password hashes and private keys.

- To encrypt the configuration file with a simple password, select **Encrypt with password** and enter the password into the **Encryption password** and **Confirm password** fields.

Note
SSB accepts passwords that are not longer than 150 characters. The following special characters can be used: `!"#$%&'()*+,-./:;<=>?@[\]^-`{|}`

- To encrypt the configuration file with GPG, select **GPG encryption**. Note that this option uses the same GPG key that is used to encrypt automatic system backups, and is only available if you have uploaded the public part of a GPG key to SSB at **Basic Settings > Management > System backup**. For details, see *Procedure 4.7.4, Encrypting configuration backups with GPG (p. 58)*.

Step 3. Click **Export**.

Note
The exported file is a gzip-compressed archive. On Windows platforms, it can be decompressed with common archive managers such as the *free 7-Zip tool*.

The name of the exported file is `<hostname_of_SSB>-YYYYMMDDTHHMM.config`; the `-encrypted` or `-gpg` suffix is added for password-encrypted and GPG-encrypted files, respectively.

6.3.8. Procedure – Importing the configuration of SSB

Purpose:

The configuration of SSB can be imported from the **Basic Settings > System** page. Use the respective action buttons to perform the desired operation.

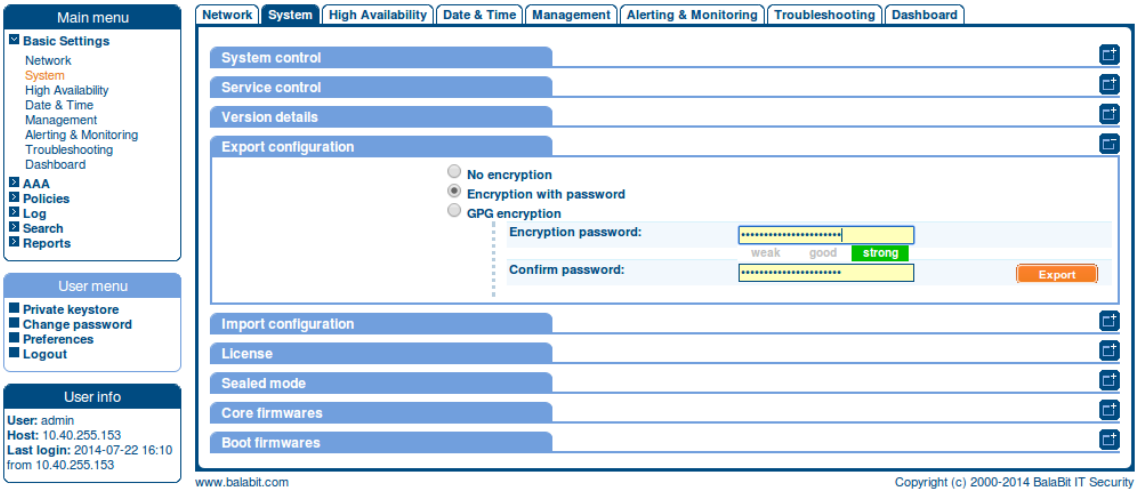

Figure 6.8. Importing the SSB configuration

Warning
It is possible to import a configuration exported from SSB 2.0 or 3.0 into SSB 4 F1, but it is not possible to restore an 1.1 or 1.0 backup into 4 F1.

Steps:

Step 1. Navigate to **Basic Settings > System > Import configuration**.

Step 2. Click **Browse** and select the configuration file to import.

Step 3. Enter the password into the **Encryption password** field and click **Upload**.

Note
SSB accepts passwords that are not longer than 150 characters. The following special characters can be used: `!"#$%&'()*+,-./:;<=>?@[\]^-`{|}`

Warning
When importing an older configuration, it is possible that there are logspaces on SSB that were created after the backing up of the old configuration. In such case, the new logspaces are not lost, but are deactivated and not configured. To make them accessible again, you have to:

1. Navigate to **Log > Spaces** and configure the logspace. Filling the **Access Control** field is especially important, otherwise the messages stored in the logspace will not be available from the **Search > Logs** interface.

2. Adjust your log path settings on the **Log > Paths** page. Here you have to re-create the log path that was sending messages to the logspace.

6.4. Accessing the SSB console

This section describes how to use the console menu of SSB, how to enable remote SSH access to SSB, and how to change the root password from the web interface.

6.4.1. Using the console menu of SSB

Connecting to the syslog-ng Store Box locally or remotely using Secure Shell (SSH) allows you to access the console menu of SSB. The console menu provides access to the most basic configuration and management settings of SSB. It is mainly used for troubleshooting purposes; the primary interface of SSB is the web interface.

Note
Detailed host information is displayed in the shell prompt:

The format of the bash prompt is:

```
(firmware_type/HA_node/hostname)username@HA_node_name:current_working_directory#
```

For example:

```
(core/master/documentation-ssb)root@ssb1:/etc#
```

- *firmware_type* is either *boot* or *core*
- *HA_node* is either *master* or *slave*
- *hostname* is the FQDN set on the GUI
- *username* is always *root*

The console menu is accessible to the *root* user using the password set during completing the Welcome Wizard.

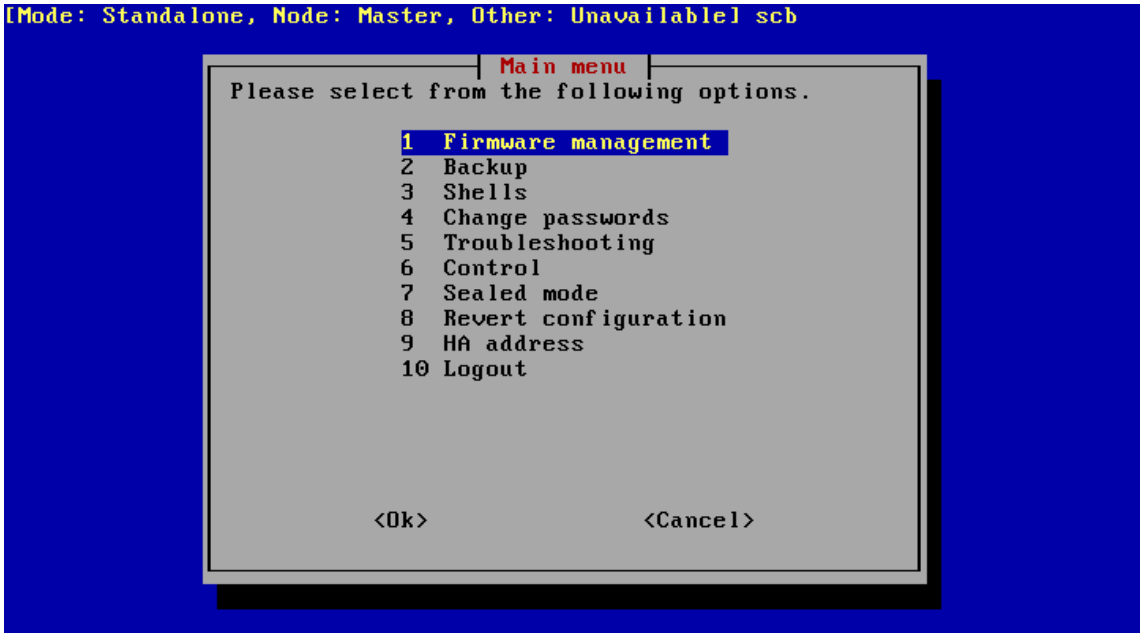

Figure 6.9. The console menu

The console menu provides allows you to perform the following actions:

- Select the active core and boot firmwares, and delete unneeded firmwares. Accessing the firmware management is useful if after an update the new firmware does not operate properly and the web interface is not available to activate the previous firmware.
- Start backup processes.
- Change the passwords of the `root` and `admin` users.
- Access the local shells of the core and boot firmwares. This is usually not recommended and only required in certain troubleshooting situations.
- Access the network-troubleshooting functions and display the available log files.
- Reboot and shutdown the system.
- Enable and disable sealed mode. For details, see *Section 6.5, Sealed mode (p. 107)*.
- Revert the configuration file. For details, see *Procedure 6.3.5, Reverting to an older firmware version (p. 98)*.
- Set the IP address of the HA interface.

Note
Note that logging in to the console menu automatically locks the SSB interface, meaning that users cannot access the web interface while the console menu is used. The console menu can be accessed only if there are no users accessing the web interface. The connection of web-interface users can be terminated to force access to the console menu.

6.4.2. Procedure – Enabling SSH access to the SSB host

Purpose:

Exclusively for troubleshooting purposes, you can access the SSB host using SSH. Completing the Welcome Wizard automatically disables SSH access. To enable it again, complete the following steps:

Warning
Accessing the SSB host directly using SSH is not recommended nor supported, except for troubleshooting purposes. In such case, the BalaBit Support Team will give you exact instructions on what to do to solve the problem.

Enabling the SSH server allows you to connect remotely to the SSB host and login using the `root` user. The password of the root user is the one you had to provide in the Welcome wizard. For details on how to change the root password from the web interface, see *Procedure 6.4.3, Changing the root password of SSB (p. 106)*

Steps:

Step 1. Navigate to **Basic Settings > Management > SSH settings**.

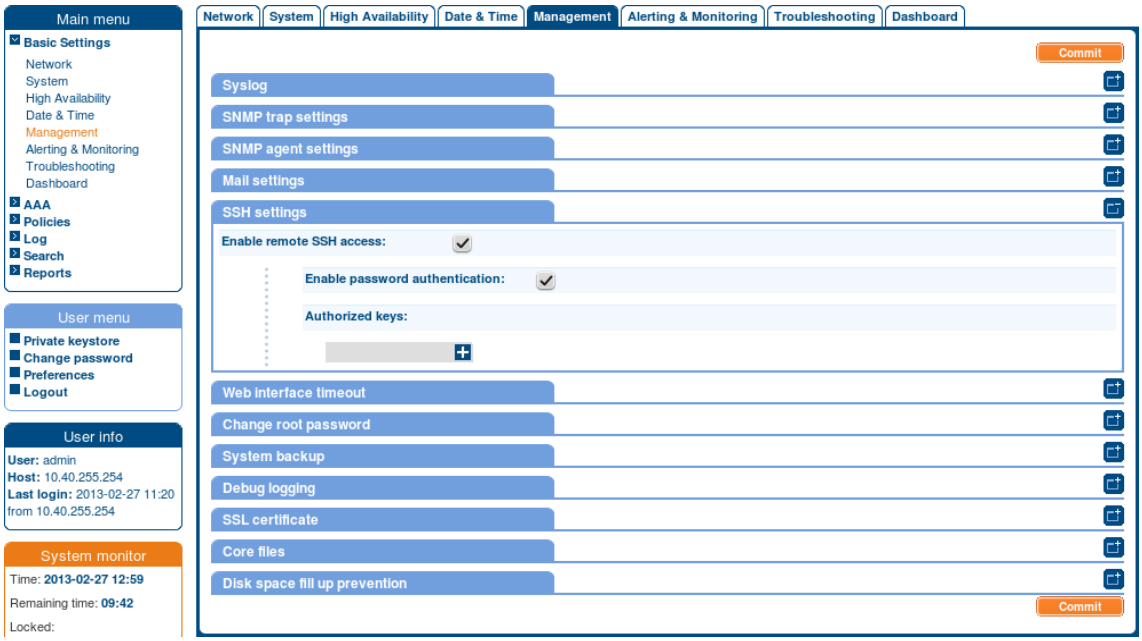

Figure 6.10. Enabling remote SSH access to SSB

Step 2. Select the **Enable remote SSH access** option.

Note
Remote SSH access is automatically disabled if Sealed mode is enabled. For details, see *Section 6.5, Sealed mode (p. 107)*.

Step 3. Set the authentication method for the remote SSH connections.

- To enable password-based authentication, select the **Enable password authentication** option.

- To enable public-key authentication, click ⊕ in the **Authorized keys** field, click ✎ and upload the private keys of the users who can access and manage SSB remotely via SSH.

Step 4. Click Commit.
The SSH server of SSB accepts connections only on the management interface if the management interface is configured. If the management interface is not configured, the SSH server accepts connections on the external interface. If possible, avoid enabling the SSH server of SSB when the management interface is not configured. For details on enabling the management connection, see *Procedure 4.3.1, Configuring the management interface (p. 37)*.

6.4.3. Procedure – Changing the root password of SSB

Purpose:

The root password is required to access SSB locally, or remotely via an SSH connection. Note that the password of the `root` user can be changed from the console menu as well. For details, see *Section 6.4, Accessing the SSB console (p. 103)*.

Steps:

Step 1. Navigate to **Basic Settings > Management > Change root password**.

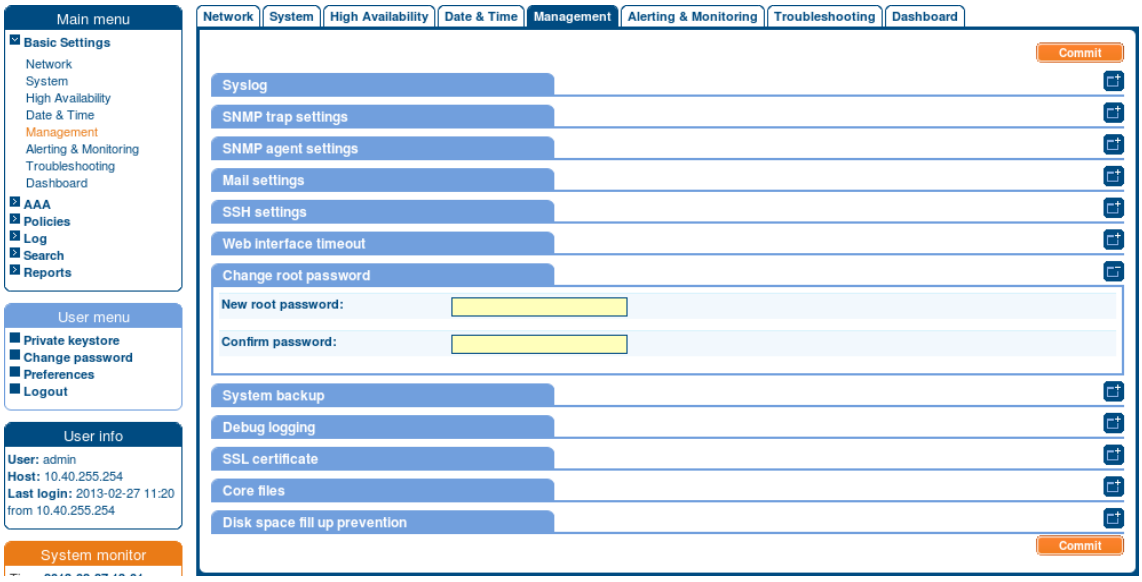

Figure 6.11. Changing the root password of SSB

Step 2. Enter the new password into the **New root password** and **Confirm password** fields.

Note
SSB accepts passwords that are not longer than 150 characters. The following special characters can be used:
`!"#$%&'()*+,-./:;<=>?@[\]^-`{|}`

Step 3. Click Commit.

6.5. Sealed mode

When sealed mode is enabled, the following settings are automatically applied:

- SSB cannot be accessed remotely via SSH for maintenance;
- the root password of SSB cannot be changed in sealed mode;
- Sealed mode can be disabled only from the local console. For details, see *Procedure 6.5.1, Disabling sealed mode (p. 107)*.

To enable sealed mode use one of the following methods:

- Select the **Sealed mode** option during the Welcome Wizard.
- Select **Basic Settings** > **System** > **Sealed mode** > **Activate sealed mode** on the SSB web interface.
- Login to SSB as root using SSH or the local console, and select **Sealed mode** > **Enable** from the console menu.

6.5.1. Procedure – Disabling sealed mode

Purpose:

To disable sealed mode, complete the following steps:

Steps:

Step 1. Go to the SSB appliance and access the local console.

Step 2. Login as `root`.

Step 3. From the console menu, select **Sealed mode** > **Disable**

Step 4. Select **Back to Main menu** > **Logout**.

6.6. Out-of-band management of SSB

SSB 4 F1 includes a dedicated out-of-band management interface conforming to the Intelligent Platform Management Interface (IPMI) v2.0 standards. The IPMI interface allows system administrators to monitor the system health of SSB and to manage the computer events remotely, independently of the operating system of SSB. SSB is accessible using the IPMI interface only if the IPMI interface is physically connected to the network.

- For details on connecting the IPMI interface, see *Procedure B.1, Installing the SSB hardware (p. 238)*.
- For details on configuring the IPMI interface, see *Procedure 6.6.1, Configuring the IPMI interface (p. 108)*.
- For details on using the IPMI interface to remotely monitor and manage SSB, see the following document:

The *Onboard BMC/IPMI User's Guide*, available on the BalaBit Hardware Documentation page at *https://www.balabit.com/support/documentation/*.

Basic information about the IPMI interface is available also on the SSB web interface on the **Basic Settings** > **High Availability** page. The following information is displayed:

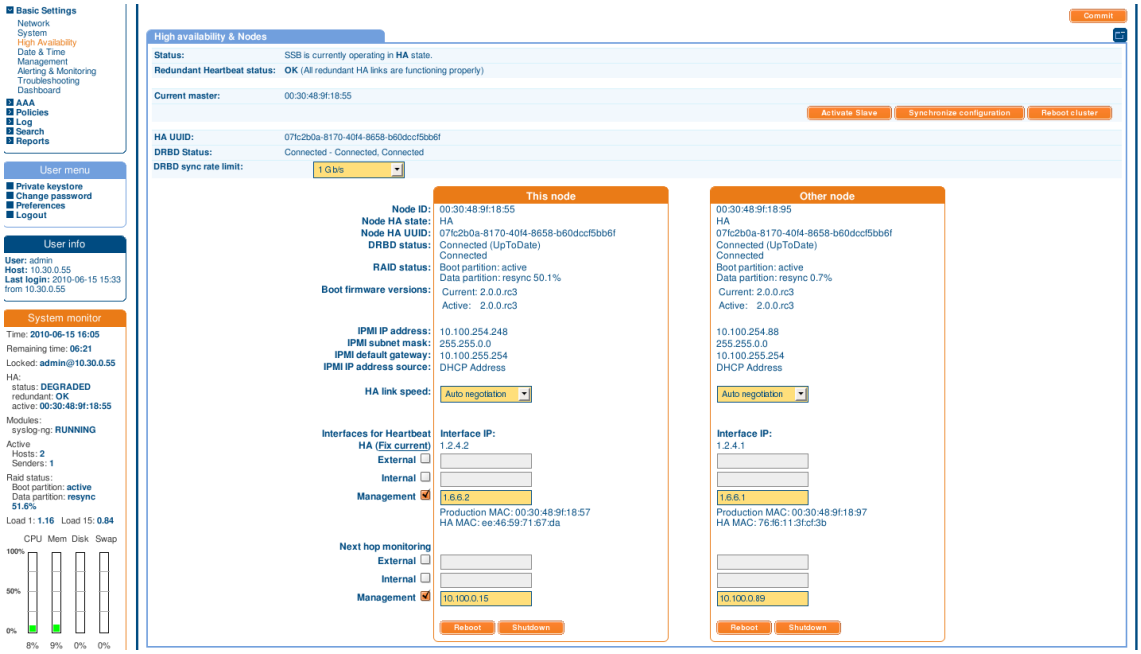

Figure 6.12. Information about the IPMI interface SSB

- **Hardware serial number**: The unique serial number of the appliance.
- **IPMI IP address**: The IP address of the IPMI interface.
- **IPMI subnet mask**: The subnet mask of the IPMI interface.
- **IPMI default gateway IP**: The address of the default gateway configured for the IPMI interface.
- **IPMI IP address source**: Shows how the IPMI interface receives its IP address: dynamically from a DHCP server, or it uses a fixed static address.

6.6.1. Procedure – Configuring the IPMI interface

Purpose:

To modify the network configuration of IPMI from the console of SSB, complete the following steps.

Prerequisites:

SCB is accessible using the IPMI interface only if the IPMI interface is physically connected to the network. For details on connecting the IPMI interface, see *Procedure B.1, Installing the SSB hardware (p. 238)*.

Warning
IPMI searches for available network interfaces during boot. Make sure that IPMI is connected to the network through the dedicated ethernet interface before SSB is powered on.

It is not necessary for the IPMI interface to be accessible from the Internet, but the administrator of SSB must be able to access it for support and troubleshooting purposes in case vendor support is needed. The following ports are used by the IMPI interface:

- Port 623 (UDP): IPMI (cannot be changed)
- Port 5123 (UDP): floppy (cannot be changed)
- Port 5901 (TCP): video display (configurable)
- Port 5900 (TCP): HID (configurable)
- Port 5120 (TCP): CD (configurable)
- Port 80 (TCP): HTTP (configurable)

Steps:

Step 1. Use the local console (or SSH) to log in to SSB as root.

Step 2. Choose **Shells** > **Boot shell**.

Step 3. Check the network configuration of the interface:
```
# ipmitool lan print
```
This guide assumes that *channel 1* is used for LAN. If your setup differs, adjust the following commands accordingly.

Step 4. Configure the interface. You can use DHCP or configure a static IP address manually.

- To use DHCP, enter the following command:
```
# ipmitool lan set 1 ipsrc dhcp
```

- To use static IP, enter the following command:
```
# ipmitool lan set 1 ipsrc static
```
Set the IP address:
```
# ipmitool lan set 1 ipaddr <IPMI-IP>
```
Set the netmask:
```
# ipmitool lan set 1 netmask <IPMI-netmask>
```
Set the IP address of the default gateway:
```
# ipmitool lan set 1 defgw ipaddr <gateway-IP>
```

Step 5. Configure IPMI to use the dedicated Ethernet interface:

- On the N1000, T1, T4, and T10 appliances, issue the following command:
  `ipmitool raw 0x30 0x70 0xc 1 0`

- On the 1000d, 5000, and 10000 appliances, issue the following command:
  `# ipmitool raw 0x30 0x70 0xc 1 1 0`

Step 6. Verify the network configuration of IPMI:
`# ipmitool lan print 1`

Use a browser to connect to the reported network address.

Step 7. Change the default password:

Step a. Log in to the IPMI web interface using the default login credentials (username: `ADMIN`, password: `ADMIN` or `changeme`, depending on your hardware).

Note
The login credentials are case sensitive.

Step b. Navigate to **Configure** > **Users**.

Step c. Select **ADMIN**, and choose **Modify User**.

Step d. Change the password, and save the changes with **Modify**.

6.7. Managing the certificates used on SSB

SSB uses a number of certificates for different tasks that can be managed from the **Basic Settings** > **Management** > **SSL certificate** menu.

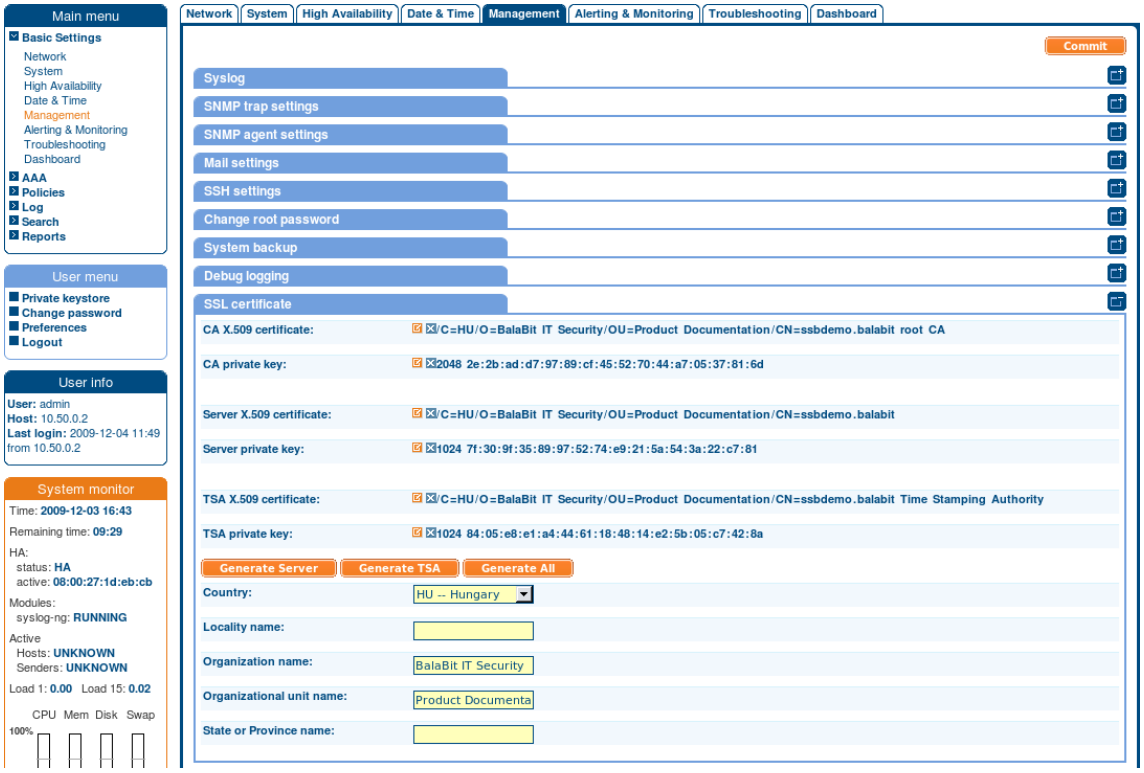

Figure 6.13. Changing the web certificate of SSB

The following certificates can be modified here:

- **CA certificate**: The certificate of the internal Certificate Authority of SSB.
- **Server certificate**: The certificate of the SSB web interface, used to encrypt the communication between SSB and the administrators.

Note
If this certificate is changed, the browser of SSB users will display a warning stating that the certificate of the site has changed.

- **TSA certificate**: The certificate of the internal Timestamping Authority that provides the timestamps used when creating encrypted logstores.

For every certificate, the distinguished name (DN) of the X.509 certificate and the fingerprint of the private key is displayed. To display the entire certificate click on the DN; to display the public part of the private key, click on the fingerprint. It is not possible to download the private key itself from the SSB web interface, but the public part of the key can be downloaded in different formats (for example PEM, DER, OpenSSH, Tectia). Also, the X.509 certificate can be downloaded in PEM and DER formats.

Note
Other parts of SSB may use additional certificates that are not managed here.

During the initial configuration, SSB creates a self-signed CA certificate, and uses this CA to issue the certificate of the web interface (see **Server certificate**) and the internal Timestamping Authority (**TSA certificate**).

There are two methods to manage certificates of SSB:

- **Recommended**: Generate certificates using your own PKI solution and upload them to SSB.

  Generate a CA certificate and two other certificates signed with this CA using your PKI solution and upload them to SSB. For the Server and TSA certificates, upload the private key as well. BalaBit recommends using 2048-bit RSA keys (or stronger).

  For details on uploading certificates and keys created with an external PKI, complete *Procedure 6.7.2, Uploading external certificates to SSB (p. 113)*.

Warning
The Server and the TSA certificates must be issued by the same Certificate Authority.

- Use the certificates generated on SSB. In case you want to generate new certificates and keys for SSB using its self-signed CA certificate, or generate a new self-signed CA certificate, complete *Procedure 6.7.1, Generating certificates for SSB (p. 112)*.

Note
Generate certificates using your own PKI solution and upload them to SSB whenever possible. Certificates generated on SSB cannot be revoked, and can become a security risk if they are somehow compromised.

6.7.1. Procedure – Generating certificates for SSB

Purpose:

Create a new certificate for the SSB webserver or the Timestamping Authority using the internal CA of SSB, or create a new, self-signed CA certificate for the internal Certificate Authority of SSB.

Steps:

Step 1. Navigate to **Basic Settings > Management > SSL certificate**.

Step 2. Fill the fields of the new certificate:

Step a. **Country**: Select the country where SSB is located (for example HU - Hungary).

Step b. **Locality**: The city where SSB is located (for example Budapest).

Step c. **Organization**: The company who owns SSB (for example Example Inc.).

Step d. **Organization unit**: The division of the company who owns SSB (for example IT Security Department).

Step e. **State or Province**: The state or province where SSB is located.

Step 3. Select the certificate you want to generate.

- To create a new certificate for the SSB web interface, select **Generate Server certificate**.
- To create a new certificate for the Timestamping Authority, select **Generate TSA certificate**.
- To create a new certificate for the internal Certificate Authority of SSB, select **Generate All**. Note that in this case new certificates are created automatically for the server and TSA certificates as well.

Note
When generating new certificates, the server and TSA certificates are signed using the certificate of the CA. If you have uploaded an external CA certificate along with its private key, it will be used to create the new server and TSA certificates. If you have uploaded an external CA certificate without its private key, use your external PKI solution to generate certificates and upload them to SSB.

Warning
Generating a new certificate automatically deletes the earlier certificate.

Step 4. Click Commit.

6.7.2. Procedure – Uploading external certificates to SSB

Purpose:

Upload a certificate generated by an external PKI system to SSB.

Prerequisites:

The certificate to upload. For the TSA and Server certificate, the private key of the certificate is needed as well. The certificates must meet the following requirements:

- SSB accepts certificates in PEM format. The DER format is currently not supported.
- SSB accepts private keys in PEM (RSA and DSA), PUTTY, and SSHCOM/Tectia format. Password-protected private keys are also supported.

Note
SSB accepts passwords that are not longer than 150 characters. The following special characters can be used: `!"#$%&'()*+,-./:;<=>?@[\]^-`{|}`

For the internal CA certificate of SSB, uploading the private key is not required.

- BalaBit recommends using 2048-bit RSA keys (or stronger).
- For the TSA certificate, the `X509v3 Extended Key Usage` attribute must be enabled and set to `critical`. Also, its default value must be set to `Time Stamping`.
- For the Server certificate, the `X509v3 Extended Key Usage` attribute must be enabled and its default value set to `TLS Web Server Authentication`. Also, the Common Name of the certificate must contain the domain name or the IP address of the SSB host.

Steps:

Step 1. Navigate to **Basic Settings** > **Management** > **SSL certificate**.

Step 2. To upload a new certificate, click next to the certificate you want to modify. A popup window is displayed.

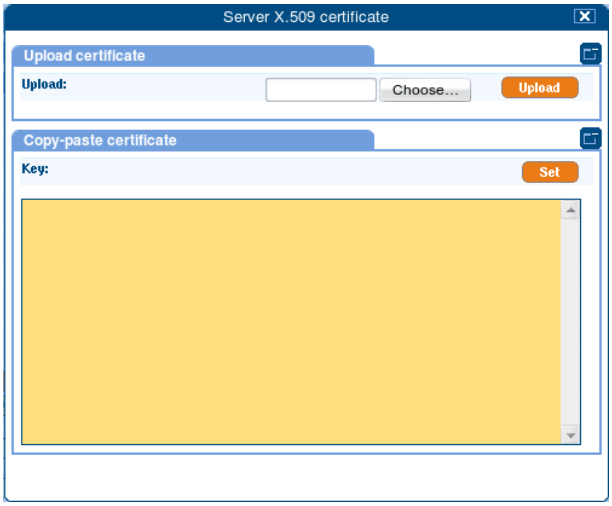

Figure 6.14. Uploading certificates

Select **Browse**, select the file containing the certificate, and click **Upload**. To upload a certificate chain, copy the certificates after each other in a single file. Alternatively, you can also copy-paste the certificate into the **Certificate** field and click **Set**. To copy-paste a certificate chain, copy and paste the certificates one by one after each other. The certificates do not have to be in order, SSB will order them. The chain is validated, if a member of the chain is missing, and error message is displayed.

Step 3. To upload the private key corresponding to the certificate, click the icon next to the private key you want to modify. A popup window is displayed.

Select **Browse**, select the file containing the certificate, and click **Upload**. Alternatively, you can also copy-paste the certificate into the **Key** field and click **Set**.

Note
In case of a certificate chain, the private key has to be the same as the bottom level certificate.

Expected result:

The new certificate is uploaded. If you receive the *`Certificate issuer mismatch`* error message after importing a certificate, you must import the CA certificate which signed the certificate as well (the private key of the CA certificate is not mandatory).

Note
To download previously uploaded certificates, click on the certificate and either download the certificate (or certificate chain) in one single PEM or DER file, or you can download single certificate files separately (if it is a certificate chain).

6.7.3. Procedure – Generating TSA certificate with Windows Certificate Authority

To generate a TSA certificate with Windows Certificate Authority (CA) that works with SSB, generate a CSR (certificate signing request) on a computer running OpenSSL and sign it with Windows CA, then import this certificate into SSB for timestamping.

Prerequisites:

A valid configuration file for OpenSSL with the following extensions:

```
[ tsa_cert ]
extendedKeyUsage = critical,timeStamping
```

You can copy `/etc/xcb/openssl-ca.cnf` from SSB to the computer that will be used for signing. Rename the file to `openssl-temp.cnf`.

Steps:

Step 1. Create CSR using the new configuration file: `openssl req -set_serial 0 -config openssl-temp.cnf -reqexts tsa_cert -new -newkey rsa:2048 -keyout timestamp.key -out timestamp.csr -nodes`

Step 2. Complete the required fields according to your environment:

```
Generating a 2048 bit RSA private key
.......................+++
.....................................+++
writing new private key to 'timestamp.key'
-----
You are about to be asked to enter information that will be incorporated
```

```
into your certificate request.
What you are about to enter is what is called a Distinguished Name or a DN.
There are quite a few fields but you can leave some blank
For some fields there will be a default value,
If you enter '.', the field will be left blank.
-----
Country Name (2 letter code) [AU]:HU
State or Province Name (full name) []:Budapest
Locality Name (eg, city) []:Budapest
Organization Name (eg, company) [Internet Widgits Pty Ltd]:BalaBit IT Security
Organizational Unit Name (eg, section) []:Service Delivery
Common Name (eg, YOUR name) []:scb35-1-i1.tohuvabohu.balabit
Email Address []:vlad@balabit.com
```

Step 3. *This step is for Windows Server 2008. Skip to the next step to continue with the instructions for Windows Server 2012.*
Sign the generated CSR with your Windows CA. Make sure that the CSR file is accessible from your Windows CA server.

Step a. To issue and sign the new certificate request, open the Microsoft Certification Authority Management Console: **Start** > **Run** and run `certsrv.msc`.

Step b. Right-click on the server name and navigate to **All Tasks** > **Submit new request...**.

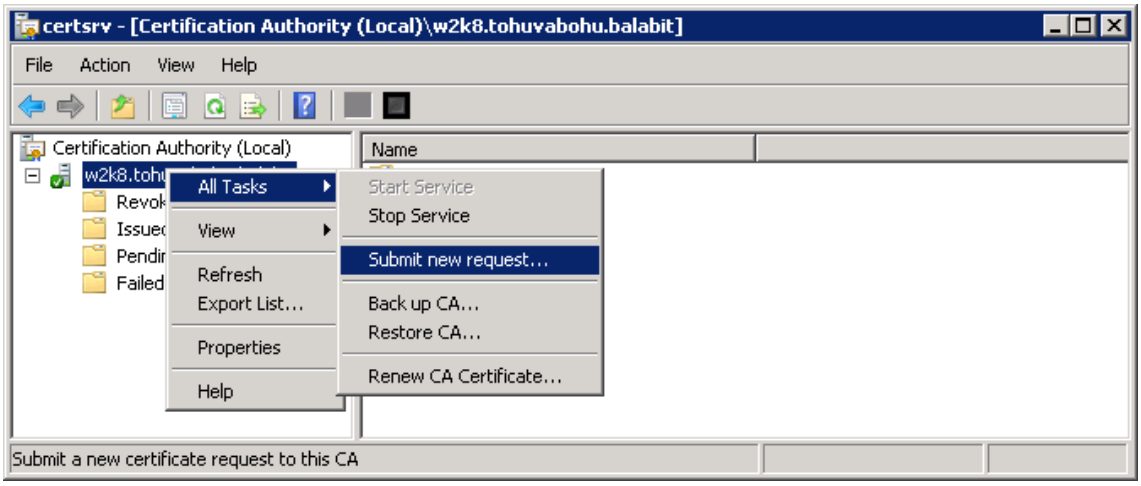

Figure 6.15. Submitting a new request

Step c. Select the CSR created in the second step.

Step d. On the left pane, click **Pending Requests**. The new certificate request is displayed in the right pane.

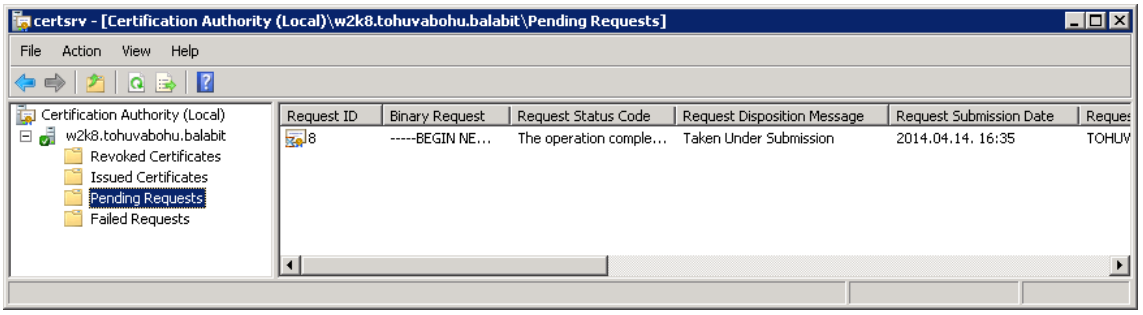

Figure 6.16. Issuing a new certificate

Step e. To issue the new SSL certificate, right-click on the pending certificate request, select “All Tasks” and click on “Issue”.

Step f. Select "Issued Certificates" and double-click on the certificate issued in the previous step.

Step g. The CA Certificate window opens. Navigate to the **Details** tab. Ensure that the required **Enhanced Key Usage** field is visible and contains the *`Time Stamping`* value.

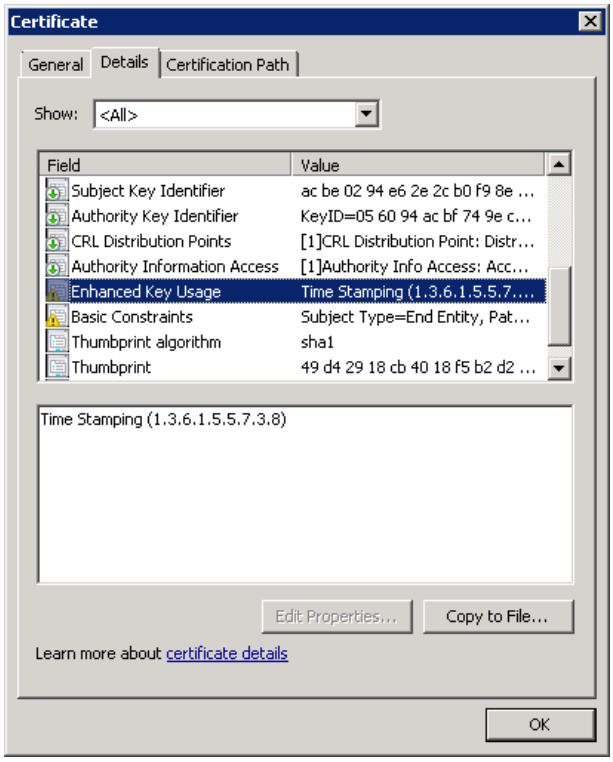

Figure 6.17. Verifying certificate details

Step h. Click **Copy to File**. The Certificate Export Wizard launches. Click **Next**.

Step i. Select the format of the certificate: **Base-64 encoded X.509 (.CER)**. Click **Next**.

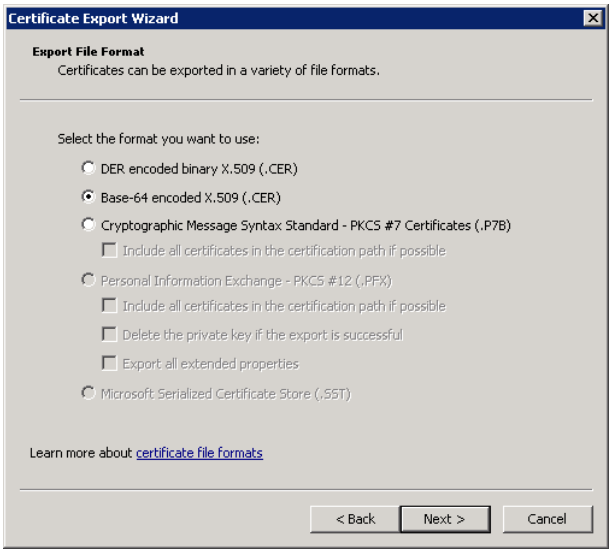
Certificate Export Wizard
Export File Format
Certificates can be exported in a variety of file formats.
Select the format you want to use:
DER encoded binary X.509 (.CER)
Base-64 encoded X.509 (.CER)
Cryptographic Message Syntax Standard - PKCS #7 Certificates (.P7B)
Include all certificates in the certification path if possible
Personal Information Exchange - PKCS #12 (.PFX)
Include all certificates in the certification path if possible
Delete the private key if the export is successful
Export all extended properties
Microsoft Serialized Certificate Store (.SST)
Learn more about certificate file formats
< Back
Next >
Cancel

Figure 6.18. Selecting certificate file format

Step j. Select location to save the certificate, and save it.

Step k. The **Completing the Certificate Export Wizard** screen is displayed. Click **Finish**.

Step 4. *This step is for Windows Server 2012.*
Create and configure a time stamping web server template in the Certificate Authority, and use that to generate the TSA certificate.

Step a. Start the Certification Authority Microsoft Management Console, and select the CA server.

Step b. Right-click on **Certificate Templates**, and choose **Manage**.

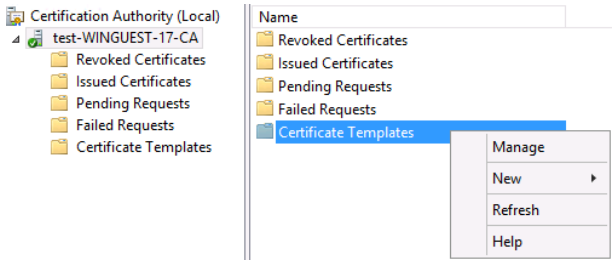

Figure 6.19. Managing certificate templates

The *Certificate Templates Console* opens.

Step c. Right-click on the **Web Server** template, and choose **Duplicate Template**.

The *Properties of New Template* window is displayed.

Step d. Make the following changes to the new template:

- On the *General* tab, change the **Template display name** to TSA.

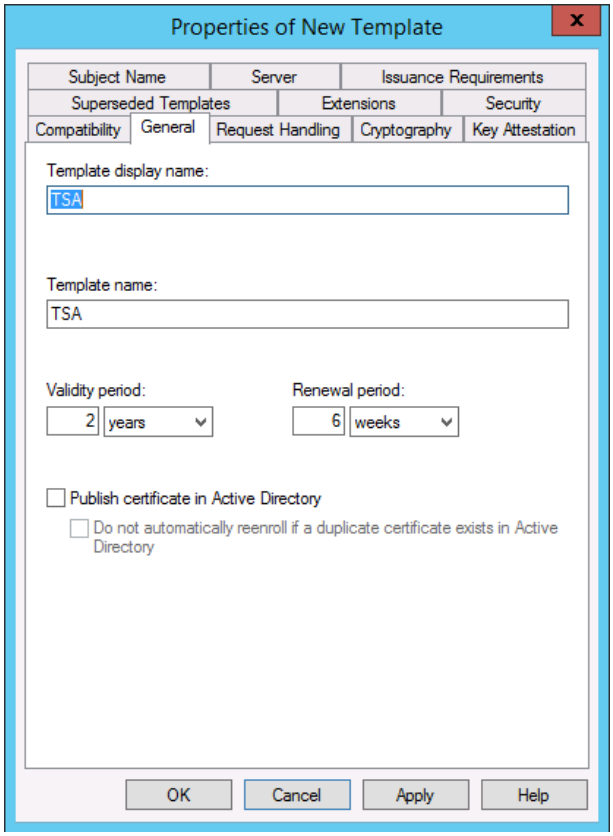

Figure 6.20. Creating the new template

- On the *Request Handling* tab, enable the **Allow private key to be exported** option.

- On the *Extensions* tab, make the following changes:

  Edit **Application Policies**: remove **Server Authentication**, add **Time Stamping**, and enable the **Make this extension critical** option, then choose **OK**.

  Edit **Key Usage**: enable the **Signature is proof of origin** option, then choose **OK**.

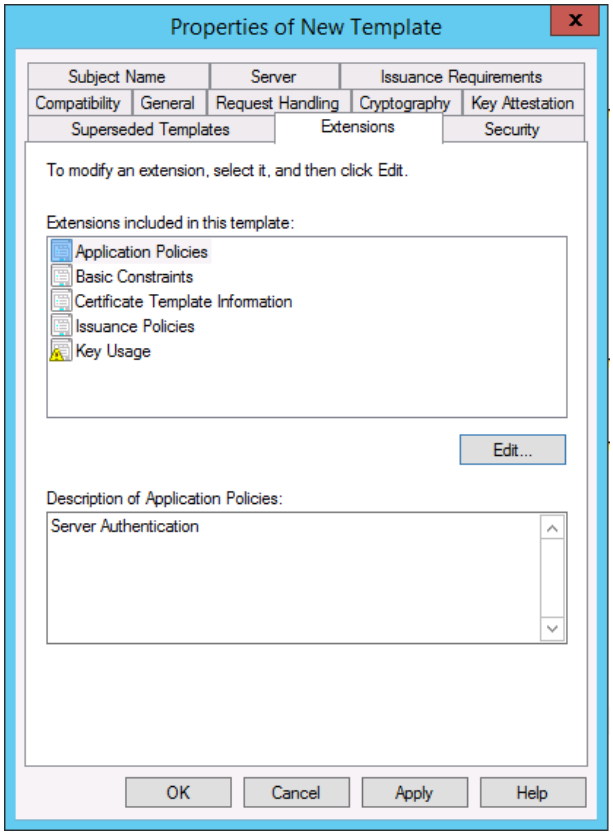

Figure 6.21. Configuring the properties of the new template

- On the *Security* tab, select **Authenticated Users**, and set **Enroll** to **Allowed**.

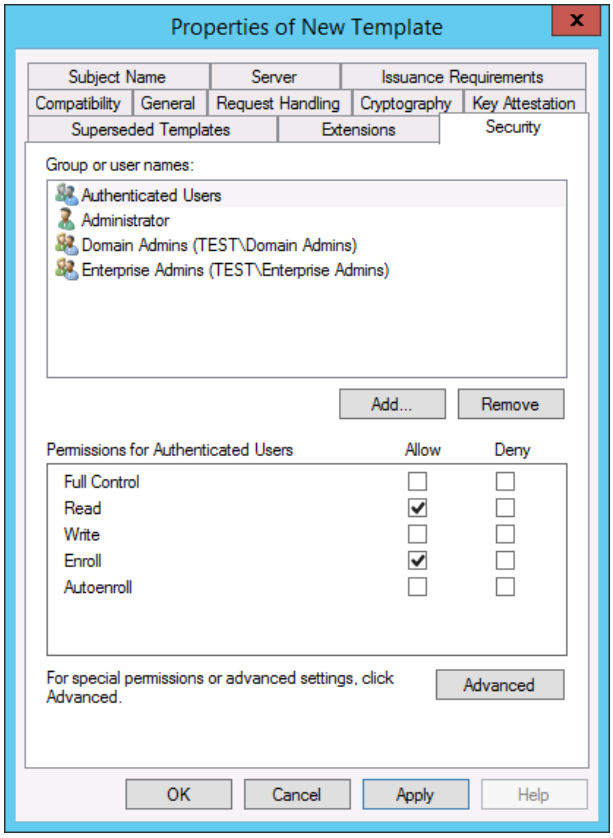

Figure 6.22. Configuring permissions for the template

Step e. Choose **Apply**. The new TSA template is now displayed in the list of templates.

Step f. Return to the Certification Authority main screen, and select the Certificate Templates folder. Right-click under the list, and choose **New > Certificate Template to Issue**. The *Enable Certificate Templates* window is displayed.

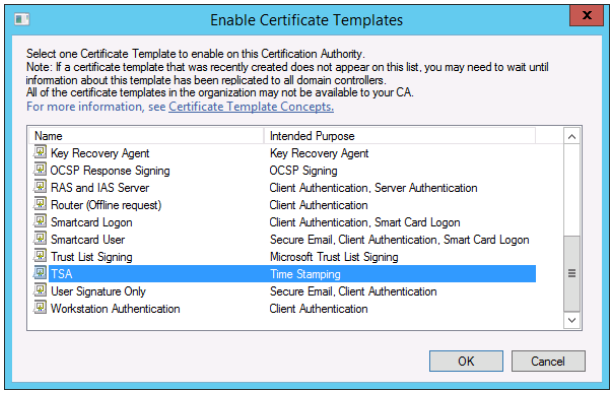
Enable Certificate Templates
Select one Certificate Template to enable on this Certification Authority.
Note: If a certificate template that was recently created does not appear on this list, you may need to wait until information about this template has been replicated to all domain controllers.
All of the certificate templates in the organization may not be available to your CA.
For more information, see Certificate Template Concepts.
Name
Intended Purpose
Key Recovery Agent
Key Recovery Agent
OCSP Response Signing
OCSP Signing
RAS and IAS Server
Client Authentication, Server Authentication
Router (Offline request)
Client Authentication
Smartcard Logon
Client Authentication, Smart Card Logon
Smartcard User
Secure Email, Client Authentication, Smart Card Logon
Trust List Signing
Microsoft Trust List Signing
TSA
Time Stamping
User Signature Only
Secure Email, Client Authentication
Workstation Authentication
Client Authentication
OK
Cancel

Figure 6.23. Enable the new template

Step g. Select the TSA certificate template, and choose **OK**.

Step h. Open the command line, and issue the following command:
```
certreq -submit -attrib "CertificateTemplate:TSA" <CSR>
```
Replace <CSR> with the full path of the CSR created earlier (in the second step).

Step i. The *Certification Authority List* is displayed. Select the CA.

Step j. The *Save Certificate* window is displayed. Choose an output folder.
The certificate is generated to the specified folder.

Step 5. In SSB, navigate to **Basic Settings** > **Management** > **SSL certificate**.

Step 6. Click next to **TSA X.509 certificate**, browse for the previously generated certificate, and click **Upload**.

Step 7. Click next to **TSA private key**, browse for the previously generated key, and click **Upload**.

Note
If the root CA (the **CA X.509 certificate** field under **Basic Settings** > **Management** > **SSL certificate**) that is used for other certificates on SSB is different from the CA that was used to sign the TSA certificate, a warning is displayed. In this scenario, ignore this warning.

6.8. Creating hostlist policies

SSB can use a list of host and network addresses at a number of places, for example for limiting the client that can send log messages to a log source, or the hosts that can access shared log spaces.

- For details on how to create a new hostlist, see *Procedure 6.8.1, Creating hostlists (p. 122)*.
- For details on how to import hostlists from a file, see *Procedure 6.8.2, Importing hostlists from files (p. 123)*.

6.8.1. Procedure – Creating hostlists

Purpose:

To create a new hostlist, complete the following steps.

Steps:

Step 1. Navigate to **Policies** > **Hostlists** and select .

Step 2. Enter a name for the hostlist (for example `servers`).

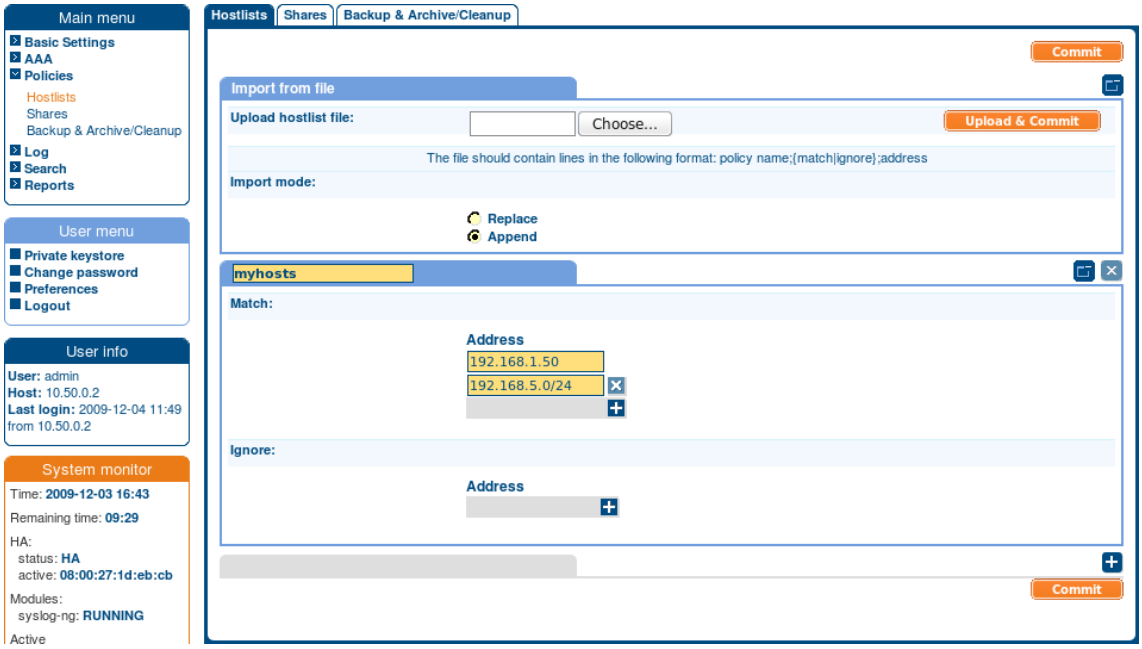

Figure 6.24. Creating hostlists

Step 3. Enter the IP address of the permitted host into the **Match** > **Address** field. You can also enter a network address in the *`IP address/netmask`* format (for example *`192.168.1.0/24`*). To add more addresses, click ✚ and repeat this step.

Step 4. To add hosts that are excluded from the list, enter the IP address of the denied host into the **Ignore** > **Address** field.

Tip
To add every address except for a few specific hosts or networks to the list, add the *`0.0.0.0/0`* network to the **Match** list, and the denied hosts or networks to the **Ignore** list.

Step 5. Click Commit.

Warning
If you modify a hostlist, navigate to **Basic Settings** > **System** > **Service control** > **Syslog traffic, indexing & search:** and select **Restart syslog-ng** for the changes to take effect.

6.8.2. Procedure – Importing hostlists from files

Purpose:

To import hostlists from a text file, complete the following steps.

Steps:

Step 1. Create a plain text file containing the hostlist policies and IP addresses to import. Every line of the file will add an IP address or network to a policy. Use the following format:
name_of_the_policy;*match* or *ignore*;*IP address*

For example, a policy that ignores the *192.168.5.5* IP address and another one that matches on the *10.70.0.0/24* subnet, use:

```
policy1;ignore;192.168.5.5
policy2;match;10.70.0.0/24
```

To add multiple addresses or subnets to the same policy, list every address or subnet in a separate line, for example:

```
policy1;ignore;192.168.7.5
policy1;ignore;192.168.5.5
policy1;match;10.70.0.0/24
```

Step 2. Navigate to **Policies** > **Hostlists** > **Import from file** > **Browse** and select the text file containing the hostlist policies to import.

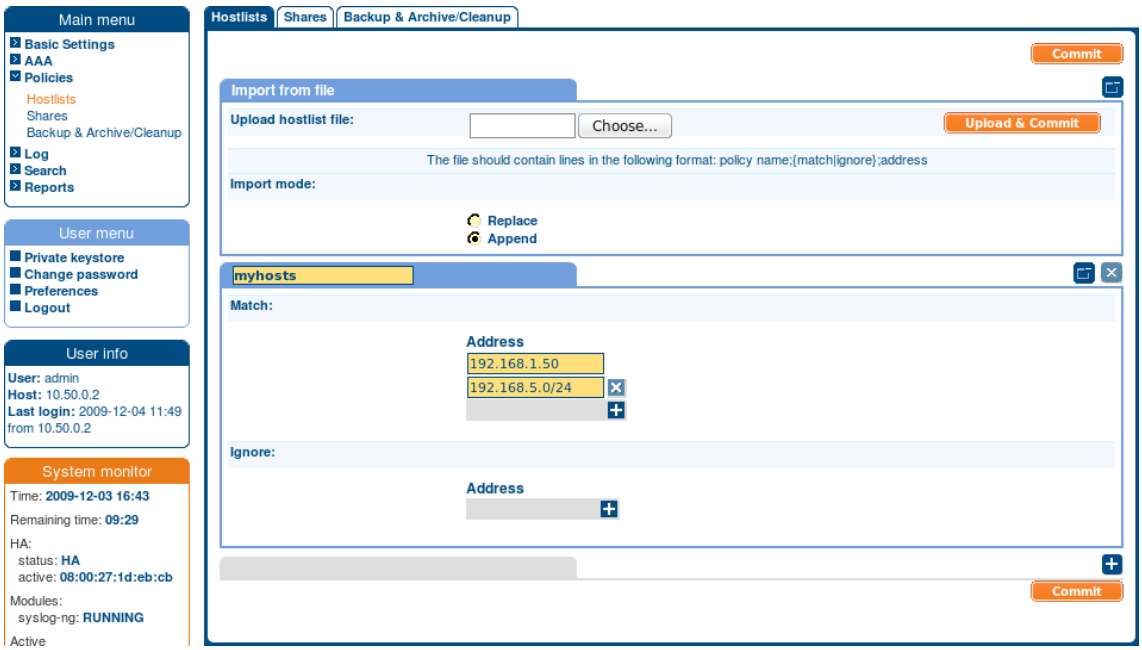

Figure 6.25. Importing hostlists

Step 3. If you are updating existing policies and want to add new addresses to them, select **Append**.

If you are updating existing policies and want to replace the existing addresses with the ones in the text file, select **Replace**.

Step 4. Click **Upload**, then Commit.

Warning
If you modify a hostlist, navigate to **Basic Settings > System > Service control > Syslog traffic, indexing & search:** and select **Restart syslog-ng** for the changes to take effect.

Chapter 7. Configuring message sources

SSB receives log messages from remote hosts via *sources*. A number of sources are available by default, but you can also create new sources. Apart from the syslog protocols, SSB can also receive messages via the SNMP protocol, and convert these messages to syslog messages.

- For details on using the built-in message sources of SSB, see *Section 7.1, Default message sources in SSB (p. 126)*.
- For details on receiving SNMP messages, see *Procedure 7.2, Receiving SNMP messages (p. 127)*.
- For details on how to create new syslog message sources, see *Procedure 7.3, Creating syslog message sources in SSB (p. 128)*.
- For details on how to create new SQL message sources, see *Section 7.4, Creating SQL message sources in SSB (p. 131)*.

7.1. Default message sources in SSB

SSB automatically accepts messages from the following built-in sources:

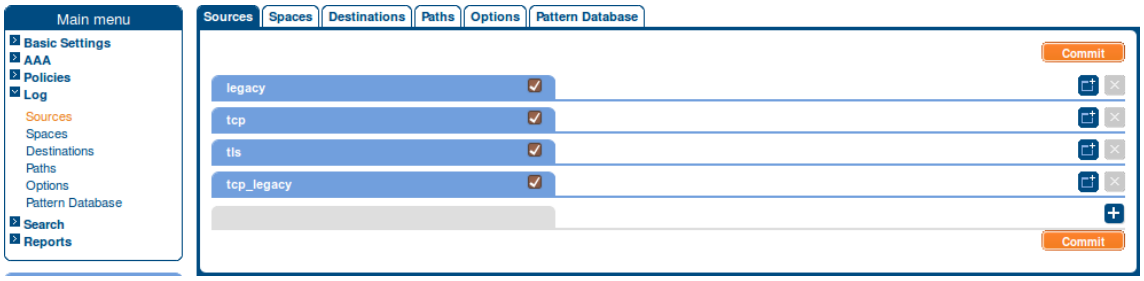

Figure 7.1. Default message sources in SSB

- *legacy*: Accepts UDP messages using the legacy BSD-syslog protocol on the port 514.
- *tcp*: Accepts TCP messages using the IETF-syslog protocol (RFC 5424) on port 601.
- *tls*: Accepts TLS-encrypted messages using the IETF-syslog protocol on port 6514. Mutual authentication is required: the client must show a (not necessarily valid) certificate; SSB sends the certificate created with the Welcome Wizard.
- *tcp_legacy*: Accepts TCP messages using the BSD-syslog protocol (RFC 3164) on port 514.

For the details of the various settings, see *Procedure 7.3, Creating syslog message sources in SSB (p. 128)*.

Note
All default sources have name resolution enabled.

7.2. Procedure – Receiving SNMP messages

Purpose:

SSB can receive SNMP messages using the SNMPv2c protocol and convert these messages to syslog messages. SNMP messages are received using a special SNMP source that can be used in log paths like any other source. To configure receiving SNMP messages, complete the following steps:

Steps:

Step 1. Navigate to **Log > Options > SNMP source**.

Step 2. Ensure that the **SNMP source** option is enabled.

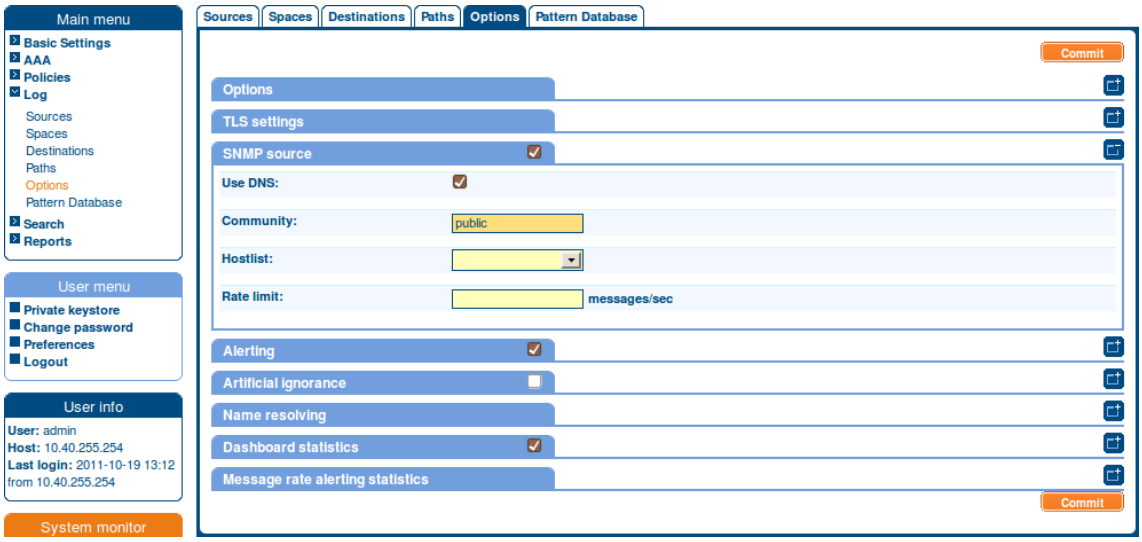

Figure 7.2. Receiving SNMP messages

Step 3. The default community of the SNMP messages is *`public`*. Modify the **Community** field if your hosts use a different community.

Note
SSB can receive messages only from a single community.

Step 4. To limit which hosts can send SNMP messages to SSB, create a hostlist policy, add the permitted hosts to the policy, and select the policy from the **Hostlist** field. For details on creating hostlists, see *Section 6.8, Creating hostlist policies (p. 122)*.

Step 5. To limit the rate of messages a host can send to SSB, enter the maximum number of packets (not messages) that SSB is allowed to accept from a single host into the **Rate limit** field. (This parameter sets the *`hashlimit`* parameter of the `iptables` packet filter that is applied to the source.)

Warning
When rate limiting is enabled, and a host sends a large number of messages, SSB processes only the amount set in the **Rate limit** field. Any additional messages are dropped, and most probably lost.

Step 6. To use name resolution for SNMP messages, enable the **Use DNS** option.

Step 7. Click Commit.

7.3. Procedure – Creating syslog message sources in SSB

Purpose:

To create a custom syslog message source, complete the following steps.

Steps:

Step 1. Navigate to **Log** > **Sources** and click ⊕.

Step 2. Enter a name for the source into the top field. Use descriptive names that help you to identify the source easily.

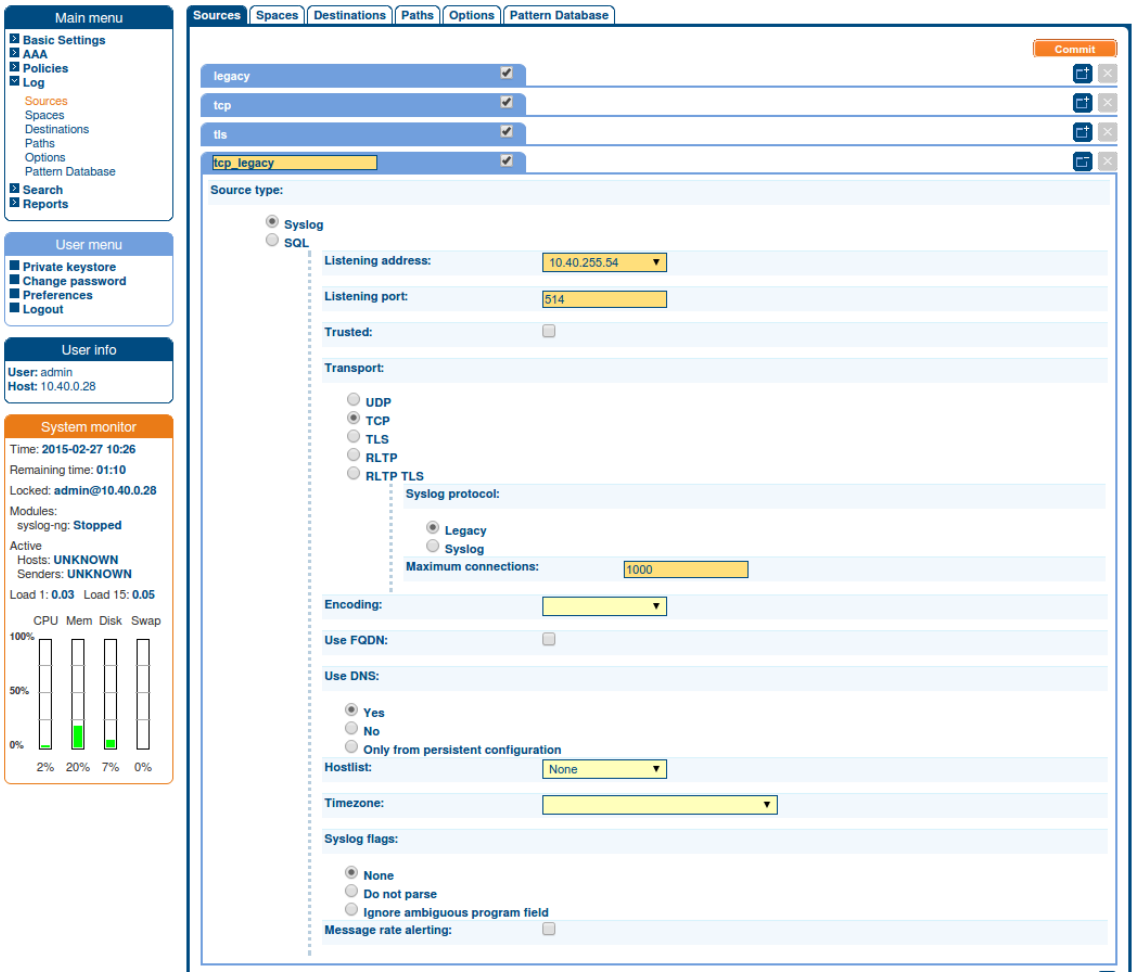

Figure 7.3. Creating new message sources

Step 3. Select **Syslog**.

Step 4. Select the interface of IP alias where SSB will receive the messages from the **Listening address** field.

Step 5. Enter the port number where SSB should accept the messages (for example *1999*).

Step 6. If the information sent by the hosts to this source can be trusted, enable the **Trusted** option. SSB keeps the timestamps and the hostname of the messages sent by trusted clients. This corresponds to enabling the *keep_timestamp()* and *keep_hostname()* syslog-ng options for the source.

Step 7. In the **Transport** field, select the networking protocol (**UDP**, **TCP**, **TLS**, **RLTP** or **RLTP TLS**) that your clients use to transfer the messages to SSB.

When using TCP or TLS, you can set the maximum number of parallel connections in the **Maximum connections** field. This option corresponds to the *max_connections()* syslog-ng parameter.

When using TLS, SSB displays a certificate to the client. This certificate can be set at **Log > Options > TLS settings** (for details, see *Procedure 11.4, Setting the certificates used in TLS-encrypted log*

transport (p. 177)). Optionally, SSB can perform mutual authentication and request and verify the certificate of the remote host (peer). Select the verification method to use from the **Peer verification** field.

- *None*: Do not request a certificate from the remote host, and accept any certificate if the host sends one.
- *Optional trusted*: If the remote host sends a certificate, SSB checks if it is valid (not expired) and that the Common Name of the certificate contains the domain name or the IP address of the host. If these checks fail, SSB rejects the connection. However, SSB accepts the connection if the host does not send a certificate.
- *Optional untrusted*: Accept any certificate shown by the remote host. Note that the host must show a certificate.
- *Required trusted*: Verify the certificate of the remote host. Only valid certificates signed by a trusted certificate authority are accepted. See *Procedure 6.7.2, Uploading external certificates to SSB (p. 113)* for details on importing CA certificates. Note that the Common Name of the certificate must contain the domain name or the IP address of the host.
- *Required untrusted*: SSB requests a certificate from the remote host, and rejects the connection if no certificate is received. However, SSB accepts the connection if:
  - the certificate is not valid (expired); or
  - the Common Name of the certificate does not contain the domain name or the IP address of the host.

When using RLTP TLS, SSB only accepts Required-trusted peer verification.

Note
For details on RLTP, see *Section 2.5, Reliable Log Transfer Protocol™ (p. 8)*.

Warning
UDP is a highly unreliable protocol, when using UDP, a large number of messages may be lost without any warning. Use TCP, TLS or RLTP whenever possible.

Step 8. In case of UDP, TCP or TLS: select the syslog protocol used by the clients from the **Syslog protocol** field. The **RLTP** and **RLTP TLS** sources only work with the IETF-syslog protocol.

- If the clients use the legacy BSD-syslog protocol (RFC3164), select **Legacy**. This protocol is supported by most devices and applications capable to send syslog messages.
- If the clients use the new IETF-syslog protocol (for example the clients are syslog-ng 3.0 applications that use the `syslog` driver, or other drivers with the `flags(syslog-protocol)` option), select **Syslog**.

In case of RLTP or RLTP TLS: enter the number of maximum connections. The default value is 1000 connections. Select **Allow compression** to allow compression on level 6. Compression level cannot be changed.

Step 9. Set the character **Encoding** and **Timezone** options of the incoming messages if needed.

Step 10. Select the **Use FQDN** option if you wish to store the full domain name of the sender host.

Step 11. Select the name resolving method to use from the **Use DNS** field.

Step 12. To accept messages only from selected hosts, create a hostlist and select it in the **Hostlist** field. For details on creating hostlists, see *Section 6.8, Creating hostlist policies (p. 122)*.

Step 13.

- If the messages arriving to the source do not comply to the standard syslog message format for some reason, select the **Syslog flags** > **Do not parse messages** option. This option completely disables syslog message parsing and treats the complete log line as the MESSAGE part of a syslog message. Other information (timestamp, host, and so on) is added automatically by SSB.

- If you want to parse messages that comply to the standard syslog message format, but disable parsing for those that do not, select the **Syslog flags** > **Ignore ambiguous program field** option. This will prevent SSB from treating the first word of the log message as the program name in case of non-standard syslog messages and thus resulting in unexpected behavior, for example, polluting the statistics.

Step 14. To configure message rate alerting for the source, see *Procedure 4.6.4, Configuring message rate alerting (p. 50)*.

Step 15. Click Commit.

Note
Note that in order to actually store the messages arriving to this source, you have to include this source in a log path. For details, see *Chapter 10, Managing log paths (p. 166)*.

Step 16. *Optional step*: If you want to receive messages using the RLTP or RLTP TLS protocol, make sure that you have configured your syslog-ng clients to transfer the messages to SSB using RLTP or RLTP TPS protocol. For details, see *Reliable Log Transfer Protocol™* in *The syslog-ng Premium Edition Administrator Guide*.

7.4. Creating SQL message sources in SSB

There are many applications that natively store their log messages in SQL databases. SSB can pull messages from SQL database tables in real-time, similarly to receiving messages over the network.

7.4.1. Procedure – Fetching the SQL database

Purpose:

To configure the parameters of the SQL database that you want to use as the message source, complete the following steps.

Steps:

Step 1. Navigate to **Log** > **Sources** and click ⊞.

Step 2. Enter a name for the source into the top field. Use descriptive names that help you to identify the source easily.

Figure 7.4. Fetching the SQL database

Step 3. Select **SQL**.

Step 4. Select the **Database type** to collect log messages from.

Step 5. Enter the hostname or the IP address of the database server to collect messages from.

Step 6. Enter the port of the database server to connect to. To use the default port of the database, click **Set Default Port**.

Step 7. Enter the name and the password of the database user.

Note
SSB accepts passwords that are not longer than 150 characters. The following special characters can be used: `!"#$%&'()*+,-./:;<=>?@[\]^-`{|}`

Step 8. Enter the database to connect to.

Step 9. Click **Test connection and fetch tables**. SSB reads the tables from the database.

Note
SSB can only read table names that contain numbers, uppercase and lowercase characters, hyphen (-), underscore (_), hashtag (#), at sign (@), or the dollar sign ($). Tables with names that contain other characters, including full stop (.), cannot be monitored.

7.4.2. Procedure – Configuring message parts in Basic mode

Purpose:

To create an SQL message source with only a few clicks, complete the following steps.

Steps:

Step 1. Select **Basic mode** for simple configuration mode. For advanced configuration settings (manually creating fetch queries, and so on), see *Procedure 7.4.3, Configuring message parts in Advanced mode (p. 136)*.

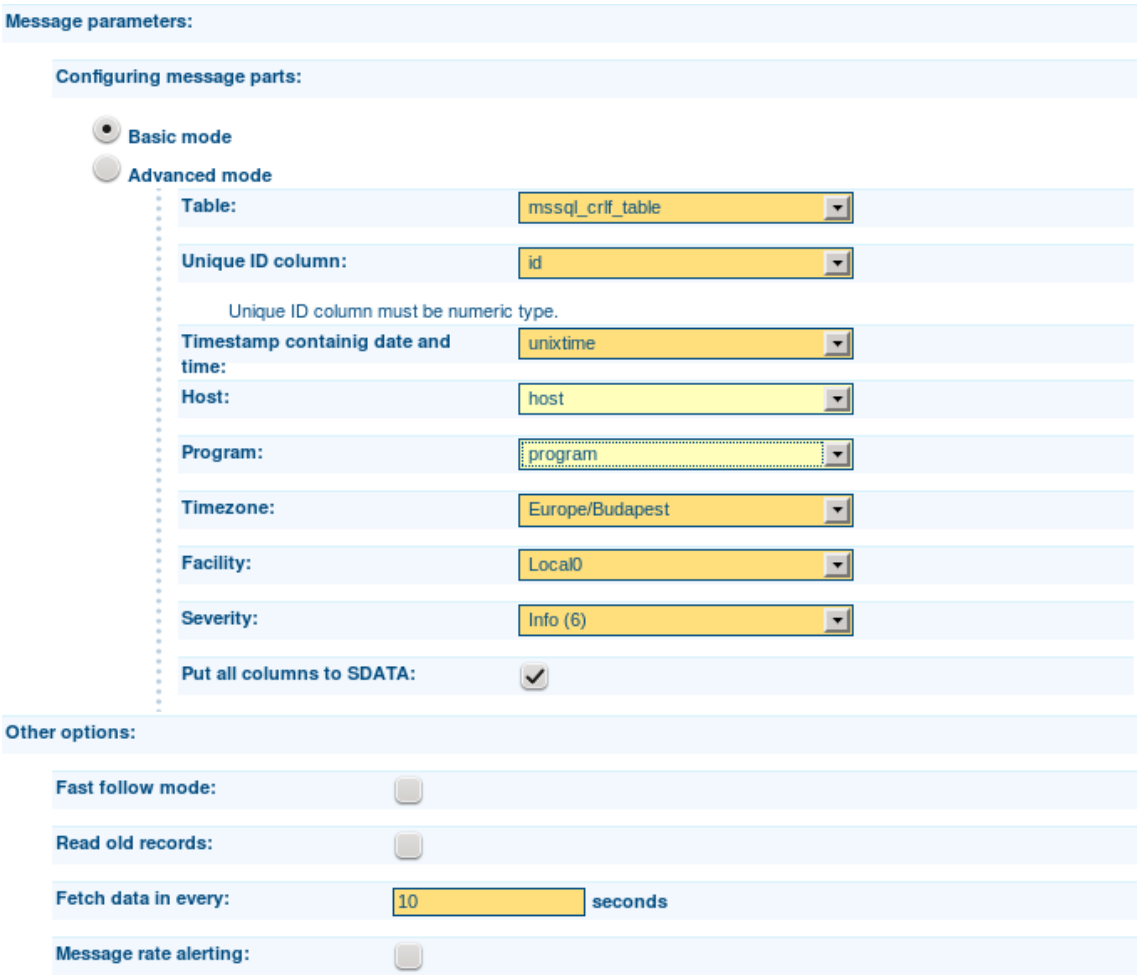

Figure 7.5. Configuring message parts in Basic mode

Step 2. Select the name of the monitored **Table**.

Note
SSB can only read table names that contain numbers, uppercase and lowercase characters, hyphen (-), underscore (_), hashtag (#), at sign (@), or the dollar sign ($). Tables with names that contain other characters, including full stop (.), cannot be monitored.

Step 3. Select the **Unique ID column**. This is the monotonically increasing unique ID of the monitored table. It must be a numeric column.

Note
SSB reads only those rows where the **Unique ID column** contains a value larger than *0*.

Step 4. Select the column containing the timestamp.

- If the timestamp column contains both date and time, select it from the list.
- If the timestamp date and timestamp time are in separate columns, select **[Set date and time separately]**. Then set the timestamp date and time columns from the respective drop-down menus.

Step 5. Optionally, select the **Host** and **Program** columns.

Step 6. Select the **Timezone**.

Step 7. Select the part of the system sending the message in **Facility**.

Step 8. Select the importance of the message in **Severity**.

Step 9. To put all columns into SDATA for further processing, enable **Put all columns into SDATA**.

Note
In Advanced mode, it is possible to put only certain selected columns (that were retrieved by the SQL query) into SDATA.

Step 10. Enable **Fast follow mode** to make syslog-ng read the database table as fast as possible.

Note
SSB reads the database periodically, each time performing one query. Each query fetches up to 3000 records. With **Fast follow mode** enabled, SSB continues querying the database until it fetched all records available at the time.

Step 11. Enable **Read old records** to make syslog-ng start reading the records from the beginning of the table, if the table has not been read yet. If it is disabled, syslog-ng will read only the new records.

Step 12. Specify the time interval between two queries by setting **Fetch data in every X seconds**. The syslog-ng application executes one query in the given timeframe (maximum 3000 records within one read operation).

Step 13. Enable **Message rate alerting** to detect abnormalities in SSB. For details, see *Procedure 4.6.4, Configuring message rate alerting (p. 50)*.

Note
In case of SQL sources, only *Messages* can be measured.

Step 14. Click **Test data retrieving**. The results are displayed in a pop-up window.

7.4.3. Procedure – Configuring message parts in Advanced mode

Purpose:

For more flexible SQL source configuration, such as manual fetch query configuration, complete the following steps.

Steps:

Step 1. Select **Advanced mode** for advanced configuration settings. For a simpler configuration, see *Procedure 7.4.2, Configuring message parts in Basic mode (p. 133)*.

Figure 7.6. Configuring message parts in Advanced mode

Step 2. Create the fetch query manually. For details, see *Section 7.4.4, Creating a fetch query manually (p. 138)*

Step 3. If you are using MSSQL database, or you encounter SQL errors or unexpected results, specify a custom query to find the last UID in the database.

Step 4. Select the **Timezone**.

Step 5. Select the part of the system sending the message in **Facility**.

Step 6. Select the importance of the message in **Severity**.

Step 7. To put all columns into SDATA for further processing, enable **Put all columns into SDATA**.

Note
In Advanced mode, it is possible to put only certain selected columns (that were retrieved by the SQL query) into SDATA.

Step 8. Enable **Fast follow mode** to make syslog-ng read the database table as fast as possible.

Note
SSB reads the database periodically, each time performing one query. Each query fetches up to 3000 records. With **Fast follow mode** enabled, SSB continues querying the database until it fetched all records available at the time.

Step 9. Enable **Read old records** to make syslog-ng start reading the records from the beginning of the table, if the table has not been read yet. If it is disabled, syslog-ng will read only the new records.

Step 10. Specify the time interval between two queries by setting **Fetch data in every X seconds**. The syslog-ng application executes one query in the given timeframe (maximum 3000 records within one read operation).

Step 11. Enable **Message rate alerting** to detect abnormalities in SSB. For details, see *Procedure 4.6.4, Configuring message rate alerting (p. 50)*.

Note
In case of SQL sources, only *Messages* can be measured.

Step 12. Click **Test data retrieving**. The results are displayed in a pop-up window.

7.4.4. Creating a fetch query manually

To create a fetch query, complete the following steps.

Warning
The SSB application does not validate or limit the contents of customized queries. Consequently, queries performed with a user with write-access can potentially modify or even harm the database. Use customized queries with care, and only for your own responsibility.

The query must return message parts with the following column names:

- *uid*:
  The uid column must contain a unique number. This number must increase monotonously. SSB will store the last read uid in $last_read_uid macro. To prevent rereading the whole table, filter records that are newer than the last read record by adding *WHERE <column_name_containing_the_id> > $last_read_uid* clause to the query. (Note that *$last_read_uid* will be substituted by SSB appropriately.)

  Add the clause *ORDER BY <column_name_containing_the_id>* at the end of the query to prevent redundant search results. .

- *datetime* or *date* and *time*:
  SSB will use the content of the datetime column as the timestamp of the log message. The following column types are supported:

  - *MySQL*: timestamp, datetime, int
  - *PostgreSQL*: timestamp, int
  - *Oracle*: timestamp, int
  - *MSSQL*: datetime, int

  If the type is int, SSB will assume that it contains a UNIX timestamp.

  When using separate date and time columns, the date column must be *date* type, the time column must be *time* type.

- *message*:
  The message field must contain the message to be logged.

- *host (optional)*:
- *program (optional)*:

The *host*, *program*, and *timezone* parameters can be selected from columns or set as a fix value. The *timezone* must contain time-shifting value and not the name of the time zone. For example, select "myhost" as *host*, "myprogram" as *program*, "+01:00" as *timezone* <further-parts-of-the-query>

Note
The query must not contain any comments.

Example 7.1. SQL source fetch_query
The following queries records that are older than the last read record:

```
SELECT * FROM <table_name> WHERE uid > $last_read_uid ORDER BY uid LIMIT 3000
```

Query to fetch the last UID from the table.

If you are using MSSQL database, or you encounter SQL errors or unexpected results, specify a custom query to find the last UID in the database.

The last UID of the table is necessary for finding the initial position in the database. By default, SSB will use the maximum value of the "uid" column from the query specified above for this purpose. However, if it does not seem to produce the required results, you can specify a custom query here

If the "Read old records" option is enabled for this database source, this field is not used.

Example 7.2. Query to fetch the last UID from the table
The following queries the last UID of the table:

```
SELECT max uid FROM <further-parts-of-the-query>
```

Note
If you are using MSSQL or MySQL database, you also have to limit the number of results of the fetch query. for example: *`SELECT top x <further-parts-of-the-query>`*. This limit must be lower than the internal SSB limit, that is 3000. In case you set a limit larger than 3000, it will be ignored and can result in performance issues.

Chapter 8. Storing messages on SSB

SSB stores log messages in binary or plain-text log files called (log) spaces. These local destinations correspond to the `logstore()` and `file()` destinations of syslog-ng. Log spaces are stored locally on the hard disk of SSB.

- For details on which logspaces are created by default, see *Section 8.1, Default logspaces in SSB (p. 140)*.
- For notes and other important information on using encrypted log files (logstores), see *Section 8.3, Using logstores (p. 141)*.
- For details on how to create additional log spaces, see *Section 8.4, Creating custom message spaces in SSB (p. 143)*.
- For details on managing logspaces, see *Section 8.5, Managing log spaces (p. 149)*.
- For details on how to make the log files accessible remotely as a network drive, see *Section 8.6, Accessing log files across the network (p. 150)*.

8.1. Default logspaces in SSB

SSB has the following log spaces by default. Any incoming message is stored in these logspaces.

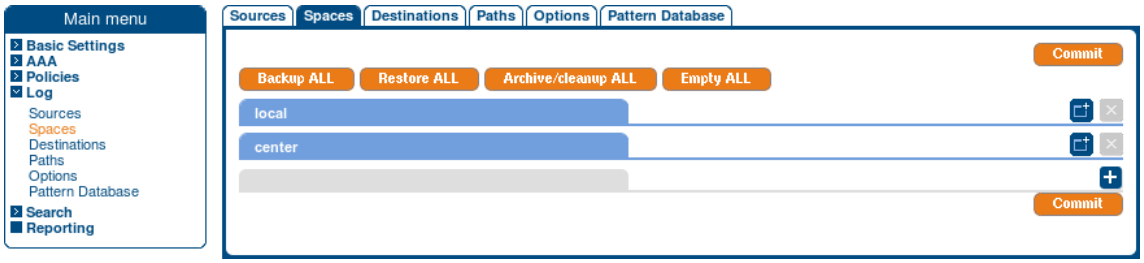

Figure 8.1. Default logspaces in SSB

- *local*: An unencrypted, binary logspace for storing the log messages of SSB.
- *center*: An unencrypted, binary logspace for storing the log messages sent by the clients.

8.2. Configuring the indexer

Navigate to **Logs** > **Spaces** and select the desired logspace.

The indexer saves the indexes for the fields that are selected and makes them searchable. Indexing fields consumes disk space and processing power.

Enter the maximum amount of memory the indexer can use for the current logspace in the **Memory limit** field.

Select the desired fields to be indexed in the **Indexed fields**. The following fields can be indexed: **Facility**, **Priority**, **Program**, **Pid**, **Host**, **Tags**, **Name/value pairs**, **Message**. For the **Name/value pairs** field, select **All** to index all Name/value fields or enter the names to be indexed in the **Only with the name** field as

comma-separated names. If the indexing of the **Message** field is enabled, the current **Delimiters** are displayed. By default, all indexers are selected.

Note
At least one field must be selected.

Note
It is not possible to search for whitespace ( ) character in the MESSAGE part of the log message, since it is a hard-coded delimiter character.

8.2.1. Limitations of the indexer

- Messages are tokenized based on the specified separator characters. Only the first 512 tokens are indexed in a message, the rest are ignored. This limitation does not affect other static fields (PROGRAM, HOST, and so on) or name-value pairs added by the pattern database or values coming from the SDATA part of incoming messages.
- Whitespaces characters (space, tabulator and so on) are always treated as delimiters.
- Tokens that are shorter than 2 characters are not indexed.
- Tokens are truncated to 59 characters. Therefore, tokens with at least 59 characters long common prefix will be handled as identical ones.
- When indexing name-value pairs, the 59 characters limitation is applied to this format: "<name-of-nvpair>=<value-of-nvpair>". Do not use long name parts, in order to avoid the premature truncation of the value part.
- The number of indexed logs cannot be more than 4294967296 (2^32) per day per logspace. This means roughly 50.000 EPS sustained traffic. If you are planning to receive and store messages at a higher sustained rate, separate them into separate logspaces.
- The shortest timeframe of searching and creating statistics is 1 minute. Smaller interval cannot be used.
- The string '`NOT`' cannot be used as the first keyword in search expressions.
- The order of the tokens in a message is not preserved. Therefore, if one message contains '`first_token second_token`' and another message contains '`second_token first_token`' search expressions such as '`first_token second_token`' will find both messages.

8.3. Using logstores

This section contains important information about using logstore files for storing log messages, and describes the current limitations of the technology. These limitations will be addressed in future versions of SSB.

- In SSB version 1.0.x, it was not possible to browse the log messages stored in encrypted logstores from the SSB web interface. This problem has been addressed in SSB 1.1; for details, see *Section 12.2, Browsing encrypted log spaces (p. 190)*.

- Indexing logstore files is currently limited. The indexer can handle only one file from a logstore for every day (SSB automatically starts a new log file for every day, this corresponds to using the *${DAY}* macro of syslog-ng). However, if you use a filename template that separates log messages based on the sender host or application, or if you use a custom template that uses a finer time-based macro (for example *${HOUR}*), then currently only the first file for the day is indexed.
- Logstore files consist of chunks. In rare cases, if the syslog-ng application running on SSB crashes for some reason, it is possible that a chunk becomes broken: it contains log messages, but the chunk was not finished completely. However, starting with SSB version 2 F1 the syslog-ng application running on SSB processes log messages into a journal file before writing them to the logstore file. That way logstore files are consistent even during unexpected crash, avoiding losing messages.
  Similarly, if the indexer application crashes for some reason, it may be possible that some parts of a logstore file are not indexed, and therefore the messages from this part of the file do not appear in search results. This does not mean that the messages are lost. Currently it is not possible to reindex a file.

These limitations will be addressed in future versions of SSB.

8.3.1. Viewing encrypted logs with logcat

To access logstore files, you can either:

- access the logstores using a network share —: for details, see *Section 8.6, Accessing log files across the network (p. 150)*(recommended), or
- login to SSB locally or remotely using SSH.

To display the contents of a logstore file, use the `logcat` command supplied with syslog-ng, for example `logcat /var/log/messages.lgs`. To display the contents of encrypted log files, specify the private key of the certificate used to encrypt the file, for example `logcat -k private.key /var/log/messages.lgs`. The contents of the file are sent to the standard output, so it is possible to use grep and other tools to find particular log messages, for example `logcat /var/log/messages.lgs |grep 192.168.1.1`.

Every record that is stored in the logstore has a unique record ID. The `logcat` application can quickly jump to a specified record using the `-- seek` option.

For files that are in use by syslog-ng, the last chunk that is open cannot be read. Chunks are closed when their size reaches the limit set in the *chunk_size* parameter, or when the time limit set in the *chunk_time* parameter expires and no new message arrives.

When the logstore file is encrypted, a hash is also generated for every chunk to verify the integrity of the chunk. The hashes of the chunks are chained together to prevent injecting chunks into the logstore file. The encryption algorithm used is *aes128* in CBC mode, the hashing (HMAC) algorithm is *hmac-sha1*.

Warning
If the syslog-ng Premium Edition application or the computer crashes, an unclosed chunk remains at the end of the file. This chunk is marked as broken, its data stays there but is not shown by `logcat`.

8.4. Creating custom message spaces in SSB

To create a custom log space, complete one of the following procedures:

- Store the log messages in binary logstore files, complete *Procedure 8.4.1, Creating a new logstore (p. 143)*.
- Store the log messages in traditional plain-text files, complete *Procedure 8.4.2, Creating a new text logspace (p. 146)*.

8.4.1. Procedure – Creating a new logstore

Steps:

Step 1. Navigate to **Log** > **Spaces** and click ⊞.

Step 2. Enter a name for the log space into the top field. Use descriptive names that help you to identify the source easily. Note that the name of the logspace must begin with a number or a letter.

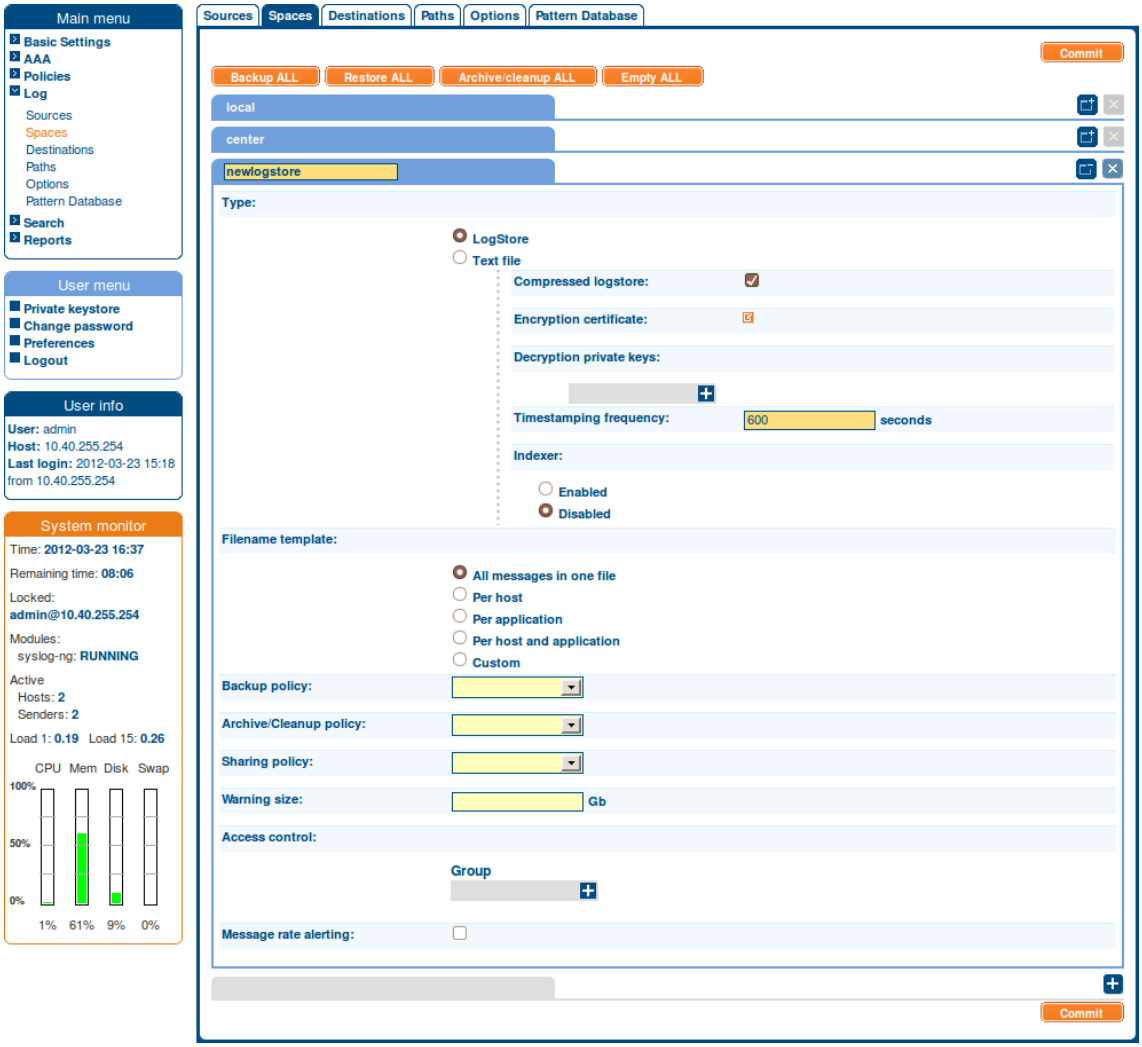

Figure 8.2. Creating a new logstore

Step 3. Select **LogStore** from the **Type** field.

Step 4. To encrypt the log files using public-key encryption, click in the **Encryption certificate** field. A popup window is displayed.

Click **Browse**, select the certificate you want to use to encrypt the log files, then click **Upload**. Alternatively, you can paste the certificate into the **Certificate** field and click **Upload**.

Note
To view encrypted log messages, you will need the private key of this certificate. For details on browsing encrypted logstores online on the SSB web interface, see *Section 12.2, Browsing encrypted log spaces (p. 190)*. Encrypted log files can be displayed using the `logcat` command-line tool as well. The `logcat` application is currently available only for UNIX-based systems.

BalaBit recommends using 2048-bit RSA keys (or stronger).

Step 5. By default, SSB requests a timestamp every ten minutes from the internal Timestamping Authority. Adjust the frequency of timestamping requests in the **Timestamping frequency** field if needed. For details on how to request timestamps from an external provider, see *Section 11.2, Timestamping configuration on SSB (p. 175)*.

Step 6. To automatically index the logstore files, select the **Enable** option of the **Indexer** field.

To limit the number of hits when searching in the logstore, enter the maximum number of search result hits in the **Maximum number of search results** field. To disable the limit, enter `0`.

By default, the following fields are indexed, if indexing is enabled: **Program**, **Host**, **Name/value pairs**, **Message**.

By default, the indexer uses the following delimiter characters to separate the message into words (tokens): `:&~?![]=,;()'"`. If your messages contain segments that include one of these delimiters, and you want to search for these segments as a whole, remove the delimiter from the list. For example, if your log messages contain MAC addresses, and you want to be able to search for messages that contain a particular MAC address, delete the colon (`:`) character from the list of delimiters. Otherwise, the indexer will separate the MAC address into several tokens.

Step 7. Logstore files are compressed by default. If you do not want to use compression, uncheck the **Compressed logstore** option.

Step 8. Select how to organize the log files of this log space from the **Filename template** field.

- To save every message received during a day into a single file, select **All messages in one file**.
- To create a separate log file for every peer (IP address or hostname) that sends messages, select the **Per host** option. This option corresponds to using the `${HOST}` macro of syslog-ng.
- To create a separate log file for every application that sends messages, select the **Per application** option. This option corresponds to using the `${PROGRAM}` macro of syslog-ng.
- To create a separate log file for every application of every peer (IP address or hostname) that sends messages, select **Per host and application** option. This option corresponds to using the `${HOST}-${PROGRAM}` macros of syslog-ng.
- To specify a custom template for naming the log files, select the **Custom** option and enter the template into the appearing **Template** field. For details on using filename templates, see *The syslog-ng Premium Edition 5 LTS Administrator Guide*.

Step 9. To create automatic daily backups of the log space to a remote server, create a backup policy and select it from the **Backup policy** field. For details on creating backup policies, see *Procedure 4.7.1, Creating a backup policy (p. 55)*.

Step 10. To archive the log space automatically daily, create an archiving policy and select it from the **Archive/Cleanup policy** field. For details on creating archiving policies, see *Procedure 4.8.1, Creating an archive policy (p. 59)*.

Warning
Use archiving and cleanup policies to remove older logfiles from SSB, otherwise the hard disk of SSB may become full.

Step 11. To make the log files of this log space available via the network, create a sharing policy and select it from the **Sharing policy** field. For details on creating sharing policies, see *Section 8.6, Accessing log files across the network (p. 150)*.

Step 12. Set a size for the log space in the **Warning size** field: SSB will send an alert if the size of this log space exceeds the limit.

Warning
Make sure that the **Logspace exceeded warning size** alert is enabled in **Basic Settings > Alerting & Monitoring** page, and that the mail and SNMP settings of the **Basic Settings > Management** page are correct. Otherwise, you will not receive any alert when the log space exceeds the size limit. For details on alerting and monitoring, see also *Section 4.6, Configuring system monitoring on SSB (p. 47)*.

Step 13. By default, members of the `search` group can view the stored messages online. Use the **Access control** option to control which usergroups can access the log space. For details, see also *Section 5.6, Managing user rights and usergroups (p. 77)*.

Step 14. Click Commit.

8.4.2. Procedure – Creating a new text logspace

Steps:

Step 1. Navigate to **Log > Spaces** and click ⊞.

Step 2. Enter a name for the log space into the top field. Use descriptive names that help you to identify the source easily.

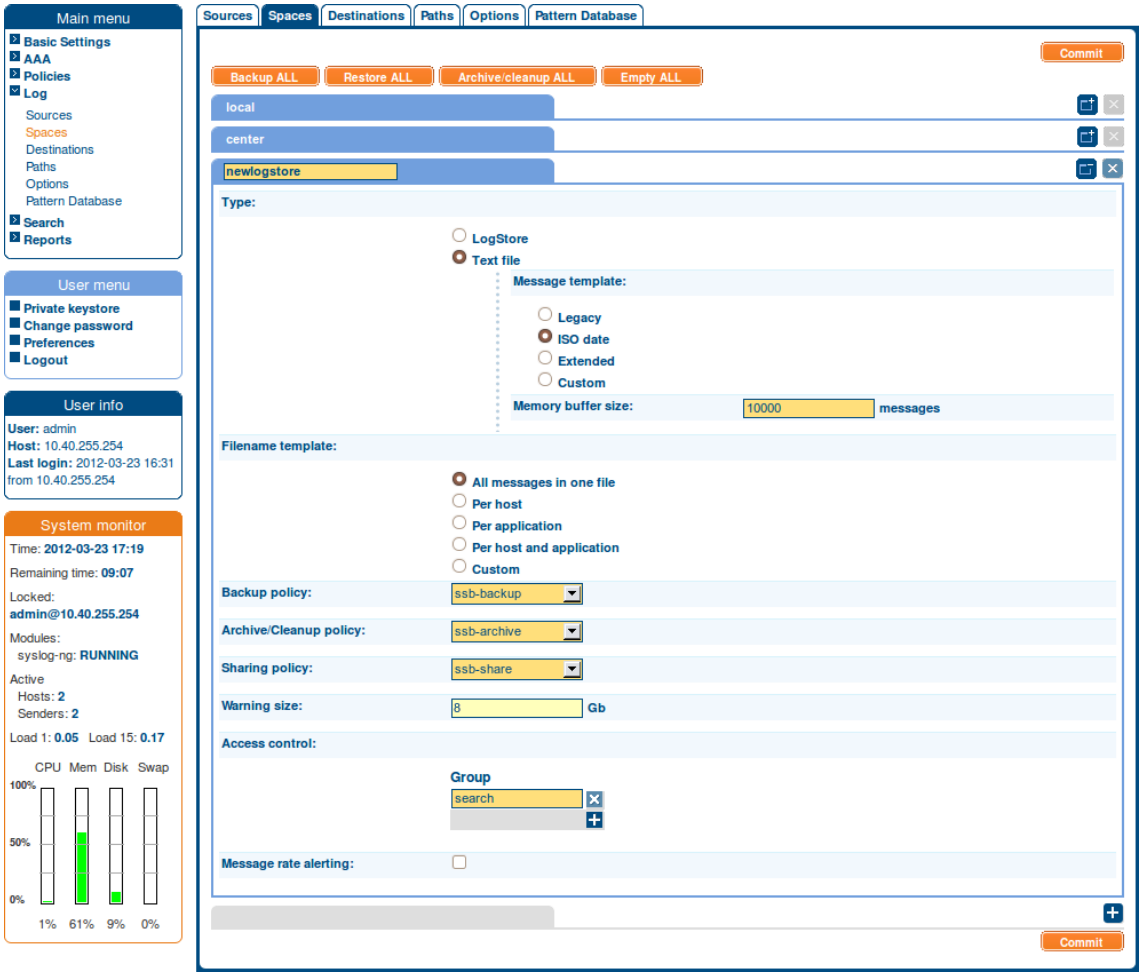

Figure 8.3. Creating a new text logspace

Step 3. Select **Text file** from the **Type** field.

Step 4. Select the template to use for the messages. The following templates are available:

- *Legacy*: `template("${DATE} ${HOST} ${MSGHDR}${MSG\n}")`
- *ISO date*: `template("${ISODATE} ${HOST} ${MSGHDR}${MSG\n}")`
- *Custom*: Specify a custom syslog-ng template in the appearing **Template** field. For details on using templates, see *The syslog-ng Premium Edition 5 LTS Administrator Guide*.

Step 5. Adjust the number of messages that are stored in the memory in the **Memory buffer size** field. This parameter corresponds to the `log_fifo_size()` parameter of syslog-ng.

Step 6. Select how to organize the log files of this log space from the **Filename template** field.

- To save every message received during a day into a single file, select **All messages in one file**.

- To create a separate log file for every peer (IP address or hostname) that sends messages, select the **Per host** option. This option corresponds to using the `${HOST}` macro of syslog-ng.
- To create a separate log file for every application that sends messages, select the **Per application** option. This option corresponds to using the `${PROGRAM}` macro of syslog-ng.
- To create a separate log file for every application of every peer (IP address or hostname) that sends messages, select **Per host and application** option. This option corresponds to using the `${HOST}-${PROGRAM}` macros of syslog-ng.
- To specify a custom template for naming the log files, select the **Custom** option and enter the template into the appearing **Template** field. For details on using filename templates, see *The syslog-ng Premium Edition 5 LTS Administrator Guide*.

Step 7. To create automatic daily backups of the log space to a remote server, create a backup policy and select it from the **Backup policy** field. For details on creating backup policies, see *Procedure 4.7.1, Creating a backup policy (p. 55)*.

Step 8. To archive the log space automatically daily, create an archiving policy and select it from the **Archive/Cleanup policy** field. For details on creating archiving policies, see *Procedure 4.8.1, Creating an archive policy (p. 59)*.

Warning
Use archiving and cleanup policies to remove older logfiles from SSB, otherwise the hard disk of SSB may become full.

Step 9. To make the log files of this log space available via the network, create a sharing policy and select it from the **Sharing policy** field. For details on creating sharing policies, see *Section 8.6, Accessing log files across the network (p. 150)*.

Step 10. Set a size for the log space in the **Warning size** field: SSB will send an alert if the size of this log space exceeds the limit.

Warning
Make sure that the **Logspace exceeded warning size** alert is enabled in **Basic Settings** > **Alerting & Monitoring** page, and that the mail and SNMP settings of the **Basic Settings** > **Management** page are correct. Otherwise, you will not receive any alert when the log space exceeds the size limit. For details on alerting and monitoring, see also *Section 4.6, Configuring system monitoring on SSB (p. 47)*.

Step 11. By default, members of the `search` group can view the stored messages online. Use the **Access control** option to control which usergroups can access the log space. For details, see also *Section 5.6, Managing user rights and usergroups (p. 77)*.

Step 12. Click Commit.

8.5. Managing log spaces

Log spaces are mostly managed automatically using backup and archiving policies, as described in *Section 4.7, Data and configuration backups (p. 55)* and *Section 4.8, Archiving and cleanup (p. 59)*. However, backup and archiving can be started manually as well. To display the details of a log space, click . A number of action buttons is shown in the top row.

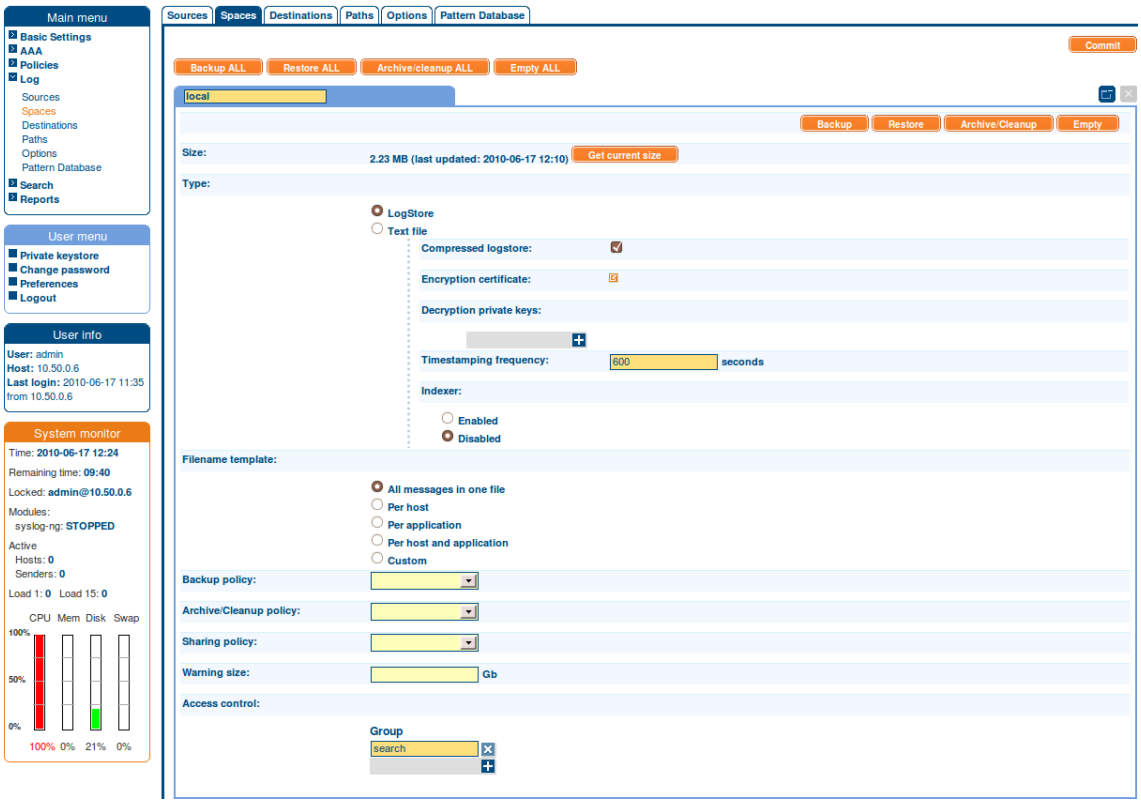

Figure 8.4. Managing log spaces

Tip
The size of the log space is displayed in the **Size** row of the log space details. To refresh the data, select **Get current size**.

- To start the backup process manually, click **Backup**.
- To restore the log files from the backup server to SSB click **Restore**.

Warning
Restoring the backup replaces every log file of the log space with the files from the backup. Any log message saved into the log space since the backup is irrevocably lost.

- To start the archiving and the cleanup process manually, click **Archive/Cleanup**.

Warning
If the archiving policy selected for the log space is set to perform only cleanup, log messages older than the Retention Time are deleted and irrevocably lost. For details, see *Section 4.8, Archiving and cleanup (p. 59)*.

- To delete every log file in the log space, click **Empty**. This option can be useful if you have to quickly free up space on SSB, or if you want to delete a log space.

Warning
This action deletes every file of the log space. Any log message not archived or backed up is irrevocably lost.

Similar action buttons are available at the top of the **Log > Spaces** page to backup, archive, or delete the contents of every logspace. These actions are performed on every logspace with their respective settings, that is, clicking **Backup All** creates a backup of every logspace using the backup policy settings of the individual logspace.

8.6. Accessing log files across the network

The log files stored on SSB can be accessed as a network share if needed using the Samba (CIFS) or Network File System (NFS) protocols. Sharing is controlled using policies that specify the type of the share and the clients (hosts) and users who can access the log files. Sharing is possible also if SSB is part of a domain.

- If you manage SSB users locally, users who have SSB account can access the shared folders. Complete *Procedure 8.6.1, Sharing log files in standalone mode (p. 150)*.
- If you manage SSB users from LDAP, you must join SSB to your domain. Complete *Procedure 8.6.2, Sharing log files in domain mode (p. 152)*.
- For details on how to access the shared files, see *Section 8.6.3, Accessing shared files (p. 155)*.

8.6.1. Procedure – Sharing log files in standalone mode

Steps:

Step 1. Navigate to **Policies > Shares > SMB/CIFS options** and select **Standalone mode**.

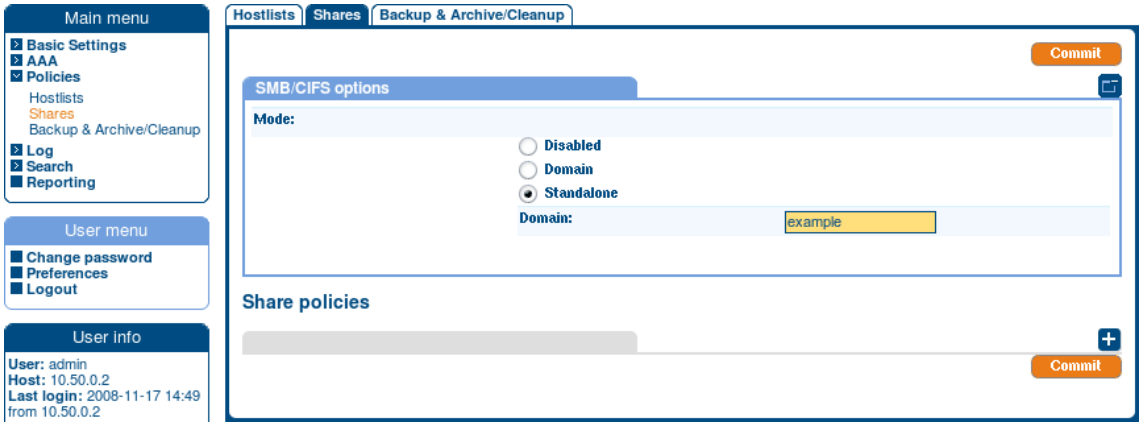

Figure 8.5. Sharing log spaces

Step 2. Select ⊞ to create a new share policy and enter a name for the policy.

Step 3. Select the type of the network share from the **Type** field.

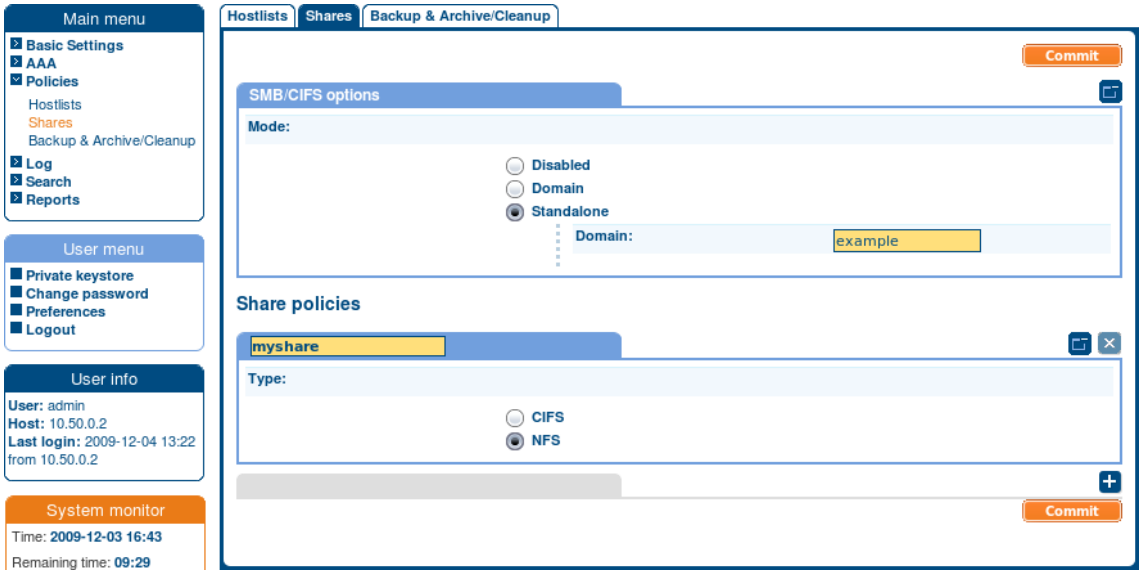

Figure 8.6. Creating share policies

- To access the log files using NFS (Network File System), select **NFS**.
- To access the log files using Samba (Server Message Block protocol), select **CIFS**.

Step 4. If you are using the Samba protocol, you can control which users and hosts can access the shares. Otherwise, every user with an SSB account has access to every shared log file.

- To control which users can access the shared files, enter the name of the usergroup who can access the files into the **Allowed group** field. For details on local user groups, see *Procedure 5.3, Managing local usergroups (p. 71)*.

- To limit the hosts from where the shares can be accessed, create a hostlist and select it from the **Hostlist** field. For details on creating hostlists, see *Section 6.8, Creating hostlist policies (p. 122)*.

Step 5. Click Commit.

Step 6. To display the details of the log space, navigate to **Log > Spaces** and click .

Step 7. Select the share policy to use from the **Sharing policy** field.

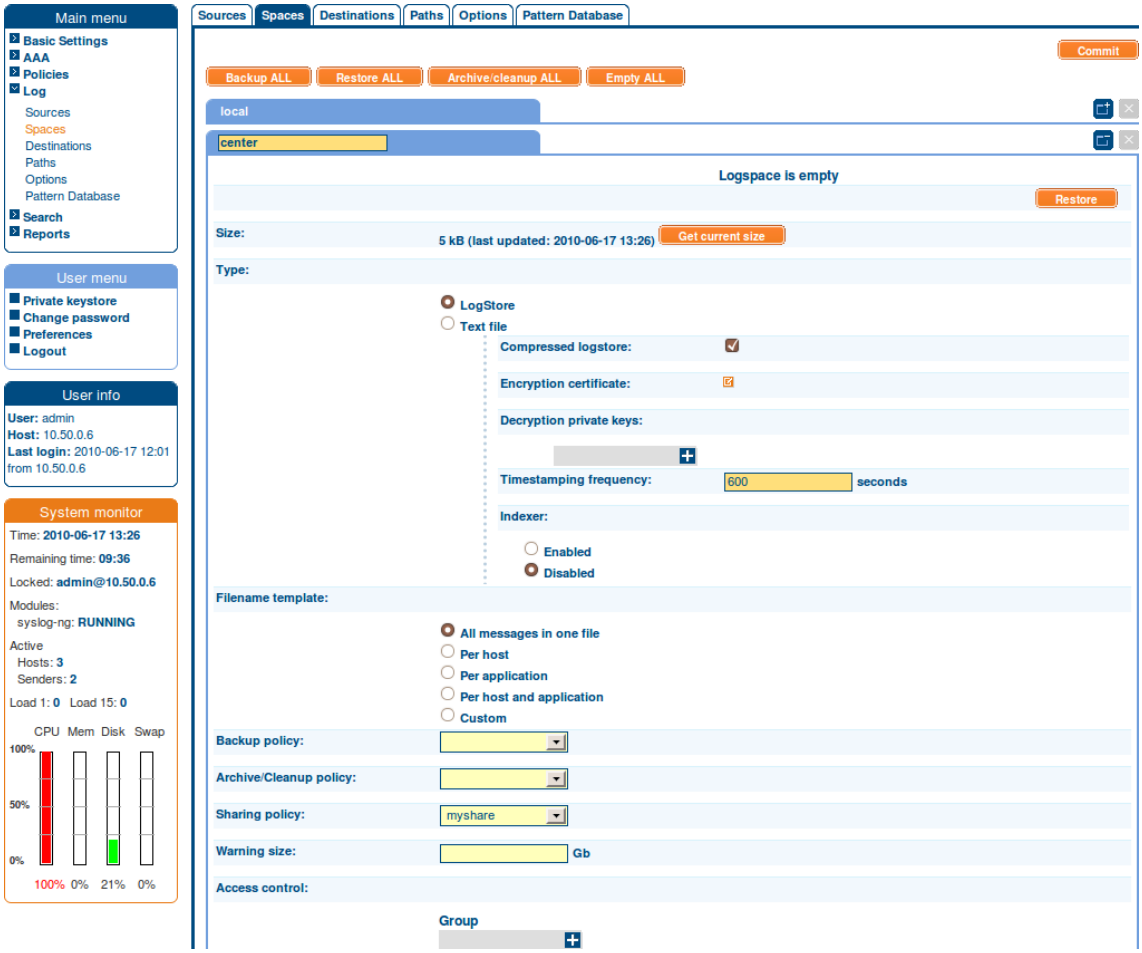

Figure 8.7. Setting the share policy of a log space

Step 8. Click Commit.

8.6.2. Procedure – Sharing log files in domain mode

Steps:

Step 1. Navigate to **Policies > Shares > SMB/CIFS options** and select **Domain mode**.

Step 2. Enter the name of the domain (for example *mydomain*) into the **Domain** field.

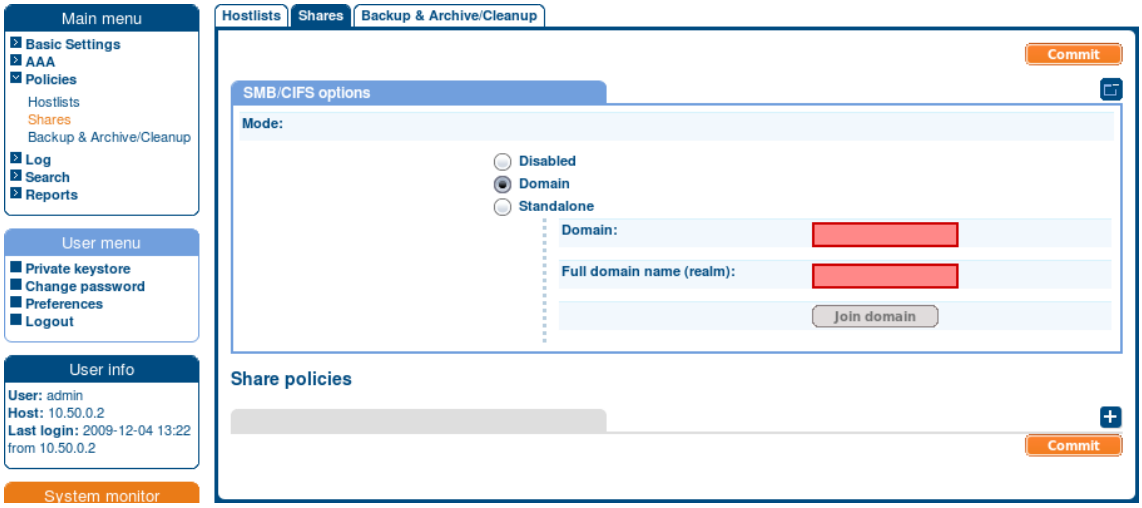

Figure 8.8. Joining a domain

Step 3. Enter the name of the realm (for example *mydomain.example.com*) into the **Full domain name** field.

Note
Ensure that your DNS settings are correct and that the full domain name can be resolved from SSB. To check this, navigate to **Basic Settings** > **Troubleshooting** > **Ping**, enter the full domain name into the **Hostname** field, and select **Ping host**.

Step 4. Click **Join domain**. A popup window is displayed.

Step 5. SSB requires an account to your domain to be able to join the domain. Enter the name of the user into the **Username** field, and the corresponding password into the **Password** field.

Note
SSB accepts passwords that are not longer than 150 characters. The following special characters can be used: *!"#$%&'()*+,-./:;<=>?@[\]^-`{|}*

Optionally, you can enter the name of your domain controller into the **Domain controller** field. If you leave this field blank, SSB will try to find the domain controller automatically.

Note
Ensure that your DNS settings are correct and that the hostname of the domain controller can be resolved from SSB. To check this, navigate to **Basic Settings** > **Troubleshooting** > **Ping**, enter the name of the domain controller into the **Hostname** field, and select **Ping host**.

Step 6. Click **Join domain**.

Step 7. Select + to create a new share policy and enter a name for the policy.

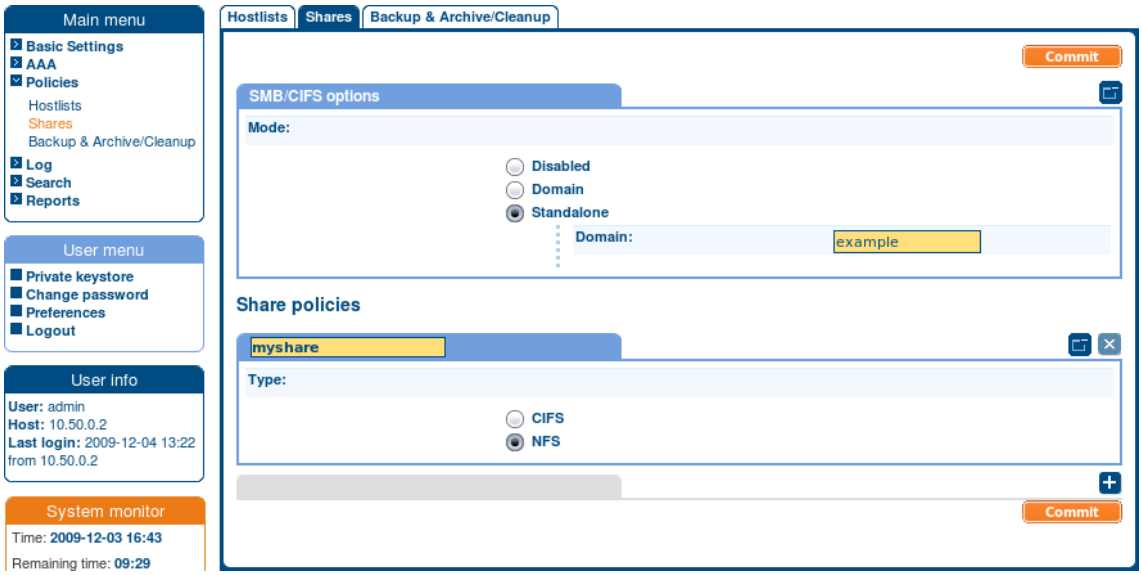

Figure 8.9. Creating share policies

Step 8. Select the type of the network share from the **Type** field.

- To access the log files using NFS (Network File System), select **NFS**.
- To access the log files using Samba (Server Message Block protocol), select **CIFS**.

Step 9. If you are using the Samba protocol, you can control which users and hosts can access the shares. Otherwise, every user with an SSB account has access to every shared log file.

- To control which users can access the shared files, enter the name of the LDAP group who can access the files into the **Allowed group** field. Note that the users and SSB must be members of the same domain.
- To limit the hosts from where the shares can be accessed, create a hostlist and select it from the **Hostlist** field. For details on creating hostlists, see *Section 6.8, Creating hostlist policies (p. 122)*.

Step 10. Click Commit.

Step 11. To display the details of the log space, navigate to **Log** > **Spaces** and click .

Step 12. Select the share policy to use from the **Sharing policy** field.

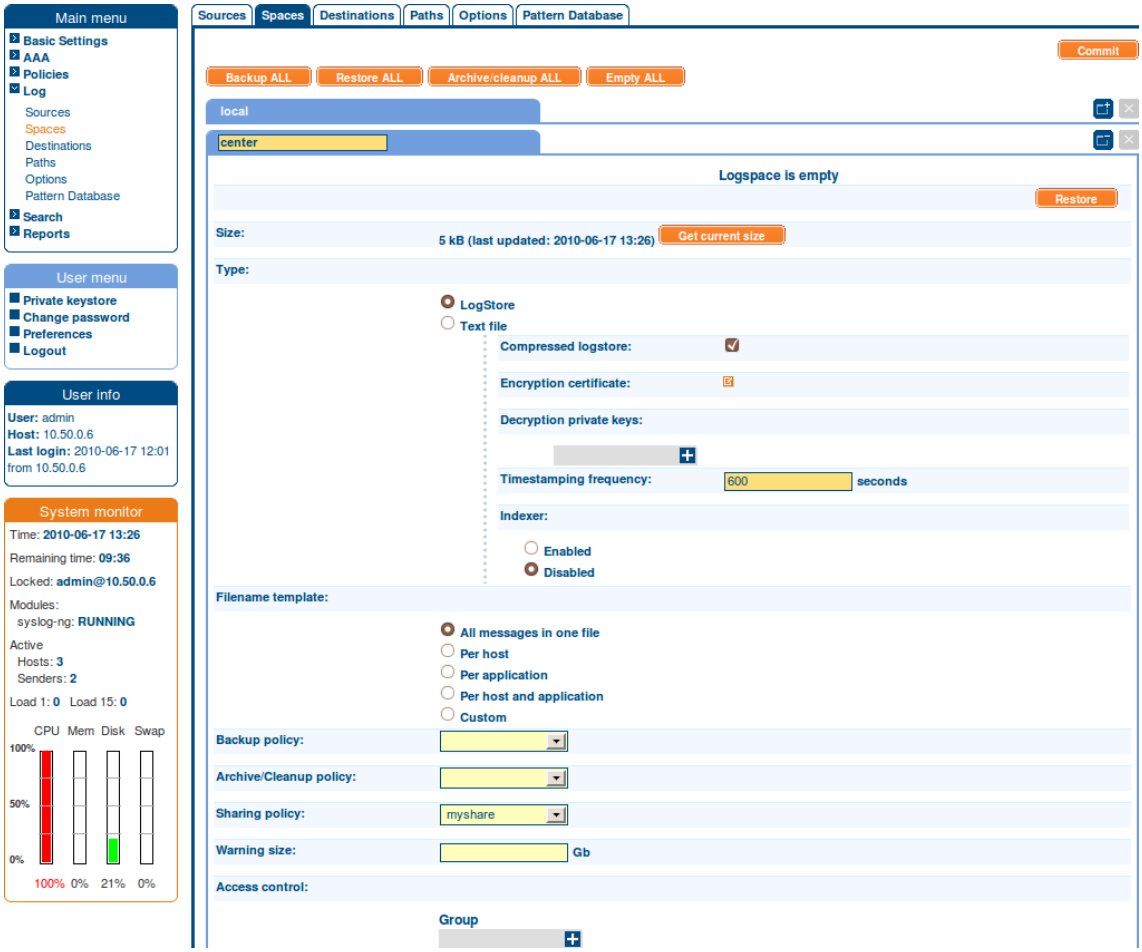

Figure 8.10. Setting the share policy of a log space

Step 13. Click Commit.

8.6.3. Accessing shared files

This section describes how to access log files that are shared using a share policy. For details on sharing log files, see *Section 8.6, Accessing log files across the network (p. 150)*.

Every shared log space is available as a separate shared folder, even if they all use a single share policy. The name of the shared folder is the name of the log space. Within the shared folder, the log files are organized into the following directory structure: `YEAR/MM-DD/`. The files are named according to the filename template set for the log space; the extension of logstore files is `.store`, while the extension of text files is `.log`. Note that the root directory of the share may also contain various files related to the log space, like index files for logstores. All files are read-only.

Note
When using NFS for sharing the log space, the name of the shared folder will be the following: `/exports/{logspace_id}/...`. The following example demonstrates how to mount a shared log space using NFS on Linux.

Example 8.1. Mounting a shared log space using NFS on Linux

```
mount -t nfs {ssb_ip}:/exports/{logspace_id} /mnt/testmount
```

Chapter 9. Forwarding messages from SSB

SSB can forward log messages to remote destinations. The remote destination can be an SQL database running on a remote server, or a syslog or log analyzing application running on a remote server.

- To forward messages to a remote SQL database, complete *Procedure 9.1, Forwarding log messages to SQL databases (p. 157)*. Currently Oracle, Microsoft SQL (MSSQL), MySQL, and PostgreSQL databases are supported.
- To forward messages to a remote server, complete *Procedure 9.3, Forwarding log messages to remote servers (p. 161)*.

9.1. Procedure – Forwarding log messages to SQL databases

Purpose:

This section describes how to forward log messages from SSB to a remote SQL database server.

Steps:

Step 1. To create a new remote destination, navigate to **Log** > **Destinations** and select .

Step 2. Enter a name for the destination.

Note
This name will be used in the name of the database tables created by SSB. For compatibility reasons, it can contain only numbers, lowercase characters, and the underscore (_) character, for example `example_database_destination`.

Step 3. Select **Database Server**.

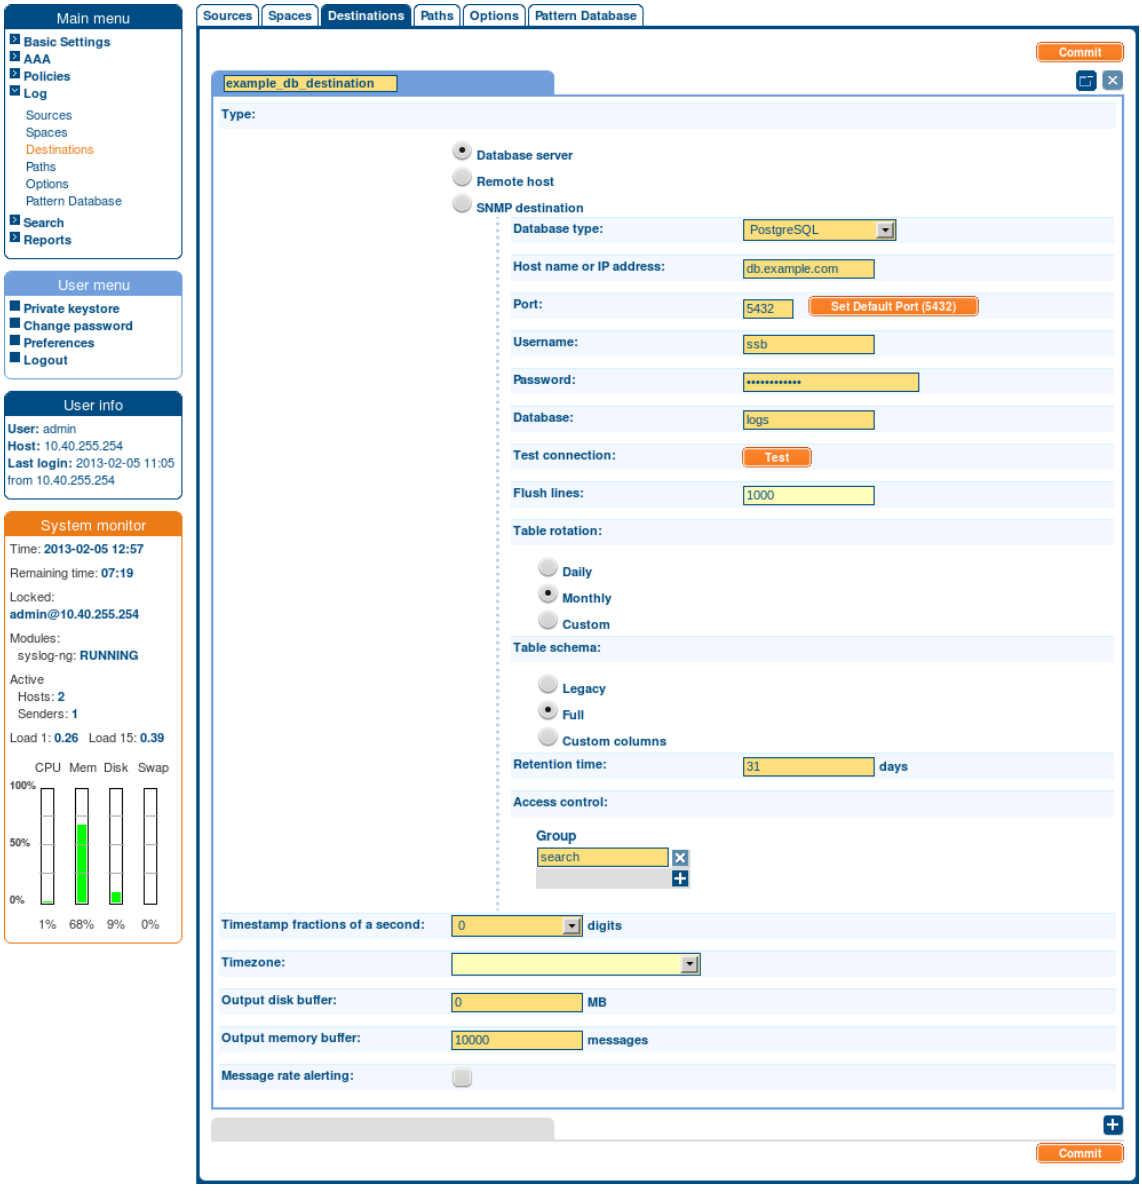

Figure 9.1. Creating database destinations

Step 4. Select the type of the remote database from the **Database type** field.

Step 5. Enter the IP address or hostname of the database server into the **Address** field. If the database is running on a non-standard port, adjust the **Port** setting.

Step 6. Enter the name and password of the database user account used to access the database into the **Username** and **Password** fields, respectively. This user needs to have the appropriate privileges for creating new tables.

Note
SSB accepts passwords that are not longer than 150 characters. The following special characters can be used: `!"#$%&'()*+,-./:;<=>?@[\]^-`{|}`

Step 7. Enter the name of the database that will store the log messages into the **Database name** field.

Step 8. *Optional step*: Enter the number of log message lines into the **Flush lines** field that SSB should wait before sending them off in a single batch. Setting this number high increases throughput as fully filled frames are sent to the network. However, it also increases message latency.

Note
Flush lines is in connection with the **Output memory buffer** value. (To set the **Output memory buffer** value, navigate to **Log > Destinations**). The value of **Output memory buffer** has to be greater than or equal to the value of **Flush lines**.

Step 9. SSB will automatically start a new table for every day or every month. Optionally, you can also create custom tables. Select the table naming template from the **Table rotation** field.

Step 10. Select which columns should SSB insert into the database. You can use one of the predefined templates, or select **Custom columns** to create a custom template. The available templates are described in *Section 9.2, SQL templates in SSB (p. 160)*.

Step 11. SSB can automatically delete older messages and tables from the database. By default, messages are deleted after one month. Adjust the **Retention time** as needed for your environment.

Step 12. The logs stored in the database can be accessed using the search interface of SSB. Enter the name of the usergroup who can access the logs into the **Access control > Group** field. To add more groups (if needed), click .

Step 13. The timestamps of most log messages is accurate only to on second. SSB can include more accurate timestamps: set how many digits should be included in the **Timestamp fractions of a second** field. This option corresponds to the `frac_digits()` parameter of syslog-ng.

Step 14. If the server and SSB are located in a different timezone and you use the `Legacy` message template (which does not include timezone information), select the timezone of the server from the **Timezone** field.

Step 15. Set the size of the disk buffer in the **Output disk buffer** field. If the remote server becomes unavailable, SSB will buffer messages to the hard disk, and continue sending the messages when the remote server becomes available. This option corresponds to the `log_disk_fifo_size()` parameter of syslog-ng.

Step 16. By default, SSB buffers up to 10000 messages in its memory if the remote server cannot accept them fast enough. To modify this value, adjust the **Output memory buffer** field as needed. This option corresponds to the `log_fifo_size()` parameter of syslog-ng.

Step 17. Click Commit.

Step 18. To start sending messages to the destination, include the new destination in a logpath. For details, see *Chapter 10, Managing log paths (p. 166)*.

Step 19. To test if the database is accessible, select **Test connection**.

9.2. SQL templates in SSB

The following sections describe the SQL templates available in SSB:

- *Legacy*
- *Full*
- *Custom*

9.2.1. The Legacy template

The **Legacy** template stores messages in the `ssb_sql_messages_${R_YEAR}_${R_MONTH}` table. The following columns are created:

- *insert_time*: The date when SSB received the message in Unixtime format.
- *rule_id*: ID of the pattern database rule that matched the message.
- *__row_id*: Identifier of the row.
- *date_time*: The date the message was sent in `YEAR-MONTH-DAY HOUR:MINUTE:SECOND` format.
- *facility*: The facility that sent the message.
- *priority*: The priority level of the message.
- *host*: The IP address or hostname of the host were the message was generated.
- *program*: The name of the application that generated the message.
- *pid*: The ID number of the process that generated the message (this field is automatically set to zero if the PID is not included in the message).
- *message*: The text of the log message.

The `insert_time`, `rule_id`, `date_time`, `facility`, `host`, and `program` columns are indexed.

9.2.2. The Full template

The **Full** template stores messages in the `ssb_sql_messages_${R_YEAR}_${R_MONTH}` table. The following columns are created:

- *insert_time*: The date when SSB received the message in Unixtime format.
- *rule_id*: ID of the pattern database rule that matched the message.
- *__row_id*: Identifier of the row.
- *date_time*: The date the message was sent in `YEAR-MONTH-DAY HOUR:MINUTE:SECOND` format.
- *facility*: The facility that sent the message.

- *priority*: The priority level of the message.
- *sourceip*: The IP address of the host that sent the message.
- *host*: The IP address or hostname of the host were the message was generated.
- *program*: The name of the application that generated the message.
- *pid*: The ID number of the process that generated the message (this field is automatically set to zero if the PID is not included in the message).
- *message*: The text of the log message.

The `insert_time`, `rule_id`, `date_time`, `facility`, `host`, `sourceip`, and `program` columns are indexed.

9.2.3. The Custom template

The **Custom** template allows you to specify the columns to use. Enter a name for the column, select its type, and specify its content using macros. For details on using macros, see *The syslog-ng Premium Edition 5 LTS Administrator Guide*. Select the **Indexed** option if you want the database to index the column.

9.3. Procedure – Forwarding log messages to remote servers

Purpose:

This section describes how to forward messages from SSB to a remote server.

Steps:

Step 1. Navigate to **Log** > **Destinations** and select ⊞ to create a new remote destination.

Step 2. Select **Remote host**.

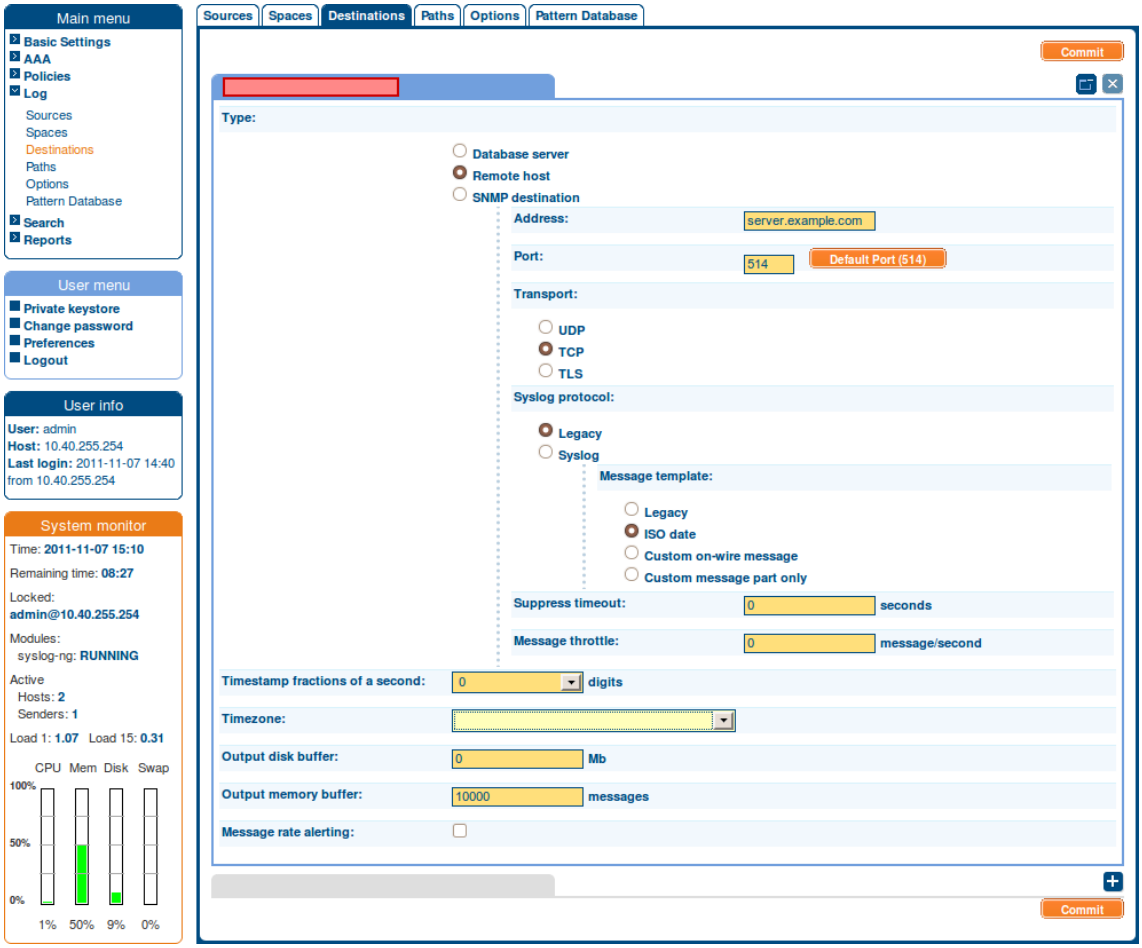

Figure 9.2. Creating server destinations

Step 3. Enter the IP address or hostname of the remote server into the **Address** field. Enter the port where the server is accepting syslog messages into the **Port** field.

Step 4. Select the network protocol used to transfer the log messages from the **Transport** field. The UDP, TCP, and the encrypted TLS protocols are available. The UDP and TLS protocols have additional parameters.

When forwarding messages using UDP, the remote host will see the messages as if they originated from SSB. Select the **Spoof source address** option to make them seem to originate from their original sender.

For TLS, select a method to verify the identity of the remote host. The following options are available:

- *None*: Do not request a certificate from the remote host, and accept any certificate if the host sends one.
- *Optional trusted*: If the remote host sends a certificate, SSB checks if it is valid (not expired) and that the Common Name of the certificate contains the domain name or the IP address of the host. If these checks fail, SSB rejects the connection. However, SSB accepts the connection if the host does not send a certificate.

- *Optional untrusted*: Accept any certificate shown by the remote host. Note that the host must show a certificate.
- *Required trusted*: Verify the certificate of the remote host. Only valid certificates signed by a trusted certificate authority are accepted. See *Procedure 6.7.2, Uploading external certificates to SSB (p. 113)* for details on importing CA certificates. Note that the Common Name of the certificate must contain the domain name or the IP address of the host.
- *Required untrusted*: SSB requests a certificate from the remote host, and rejects the connection if no certificate is received. However, SSB accepts the connection if:
  - the certificate is not valid (expired); or
  - the Common Name of the certificate does not contain the domain name or the IP address of the host.

Note
Consult the documentation of the remote server application to determine which protocols are supported.

UDP is a highly unreliable protocol and a high amount of messages may be lost without notice during the transfer. Use TCP or TLS instead whenever possible.

Step 5. Select the syslog protocol to use from the **Syslog protocol** field.

- To use the legacy BSD-syslog protocol described in RFC 3164, select **Legacy** and specify the message template to use. Select **Legacy** to use the message format described in the RFC; **ISO date** to replace the original timestamp with an ISO8061 compliant timestamp that includes year and timezone information. To customize the format of the message contents using macros, select **Custom message part only**, or **Custom on-wire message** to completely reformat the message (including the headers). For details on using macros, see *The syslog-ng Premium Edition 5 LTS Administrator Guide*. If you have no special requirements, use the **ISO date** template.

- Use the new IETF-syslog protocol. Note that most syslog applications and devices currently support only the legacy protocol. Consult the documentation of the remote server application to determine which protocols are supported. If you need, you can customize the contents of the message using macros. Note that for the IETF-syslog protocol, the header cannot be customized. For details on using macros, see *The syslog-ng Premium Edition 5 LTS Administrator Guide*.

Step 6. If SSB would send several messages with identical content to the destination, it can send only a single message and a line `Last message repeated n times.`. Enter the number of seconds to wait for identical messages into the **Suppress timeout** field. This option corresponds to the `suppress()` parameter of syslog-ng.

Step 7. To limit the maximum number of messages sent to the destination per second, enter the maximum number of messages into the **Message throttle** field. Use this output-rate-limiting functionality only when using disk-buffer as well to avoid the risk of losing messages. Specifying 0 or a lower value sets the output limit to unlimited. This option corresponds to the `throttle()` parameter of syslog-ng.

Step 8. The timestamps of most log messages is accurate only to on second. SSB can include more accurate timestamps: set how many digits should be included in the **Timestamp fractions of a second** field. This option corresponds to the *frac_digits()* parameter of syslog-ng.

Step 9. If the server and SSB are located in a different timezone and you use the *Legacy* message template (which does not include timezone information), select the timezone of the server from the **Timezone** field.

Step 10. Set the size of the disk buffer in the **Output disk buffer** field. If the remote server becomes unavailable, SSB will buffer messages to the hard disk, and continue sending the messages when the remote server becomes available. This option corresponds to the *log_disk_fifo_size()* parameter of syslog-ng.

Step 11. By default, SSB buffers up to 10000 messages in its memory if the remote server cannot accept them fast enough. To modify this value, adjust the **Output memory buffer** field as needed. This option corresponds to the *log_fifo_size()* parameter of syslog-ng.

Step 12. Click Commit.

Step 13. To start sending messages to the destination, include the new destination in a logpath. For details, see *Chapter 10, Managing log paths (p. 166)*.

9.4. Procedure – Forwarding log messages to SNMP destinations

Purpose:

To forward log messages from SSB to an SNMP destination, complete the following steps. The format of SSB SNMP messages conforms to the *CISCO-SYSLOG-MIB*.

Steps:

Step 1. Navigate to **Log** > **Destinations** and select ⊞ to create a new remote destination.

Step 2. Select **SNMP destination**.

Step 3. Enter the IP address or hostname of the SNMP destination into the **Address** field. Enter the port where the server is accepting SNMP traps into the **Port** field.

Step 4. Select the protocol version. The default value is *SNMP v2c*.

- To use the SNMP v2c protocol, select **SNMP v2c** and enter the name of the SNMP community to use in the **Community** field. The default value is *public*.

- To use the SNMP v3 protocol, select **SNMP v3**. Enter the username and the Engine ID to be used when sending SNMP traps in the respective fields. Select the authentication method to use (MD5 or SH1) and enter the authentication password. Select the encryption method to use (Disabled or DES). In case of DES, enter the encryption password.

Note
SSB accepts passwords that are not longer than 150 characters. The following special characters can be used: `!"#$%&'()*+,-./:;<=>?@[\]^-`{|}`

Step 5. The timestamps of most log messages is accurate only to on second. SSB can include more accurate timestamps: set how many digits should be included in the **Timestamp fractions of a second** field. This option corresponds to the *frac_digits()* parameter of syslog-ng.

Step 6. If the server and SSB are located in a different timezone and you use the *Legacy* message template (which does not include timezone information), select the timezone of the server from the **Timezone** field.

Step 7. Set the size of the disk buffer in the **Output disk buffer** field. If the remote server becomes unavailable, SSB will buffer messages to the hard disk, and continue sending the messages when the remote server becomes available. This option corresponds to the *log_disk_fifo_size()* parameter of syslog-ng.

Step 8. By default, SSB buffers up to 10000 messages in its memory if the remote server cannot accept them fast enough. To modify this value, adjust the **Output memory buffer** field as needed. This option corresponds to the *log_fifo_size()* parameter of syslog-ng.

Step 9. Click Commit.

Step 10. To start sending messages to the destination, include the new destination in a logpath. For details, see *Chapter 10, Managing log paths (p. 166)*.

Step 11. To properly interpret and display the SNMP messages on your destination, download and install the *CISCO-SYSLOG-MIB* in your destination software.

Chapter 10. Managing log paths

This section describes how to create and configure log paths in SSB.

- For a list of default log paths, see *Section 10.1, Default logpaths in SSB (p. 166)*.
- For details on how to create a new log path, see *Procedure 10.2, Creating new log paths (p. 166)*.
- For details on how to send only selected messages to a destination, see *Section 10.3, Filtering messages (p. 169)*.

10.1. Default logpaths in SSB

Two log paths are available by default in SSB (see **Log** > **Paths**):

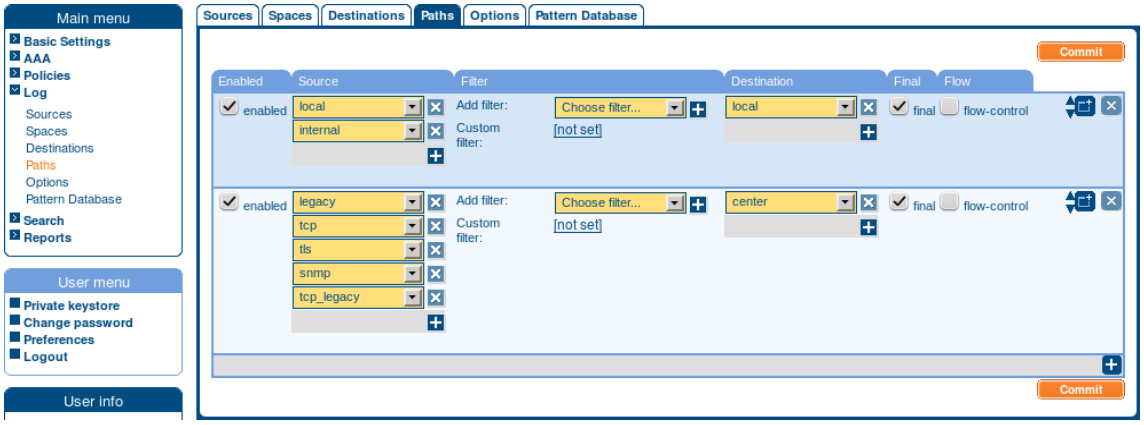

Figure 10.1. Default logpaths of SSB

- The first log path collects the local messages of SSB. It sends every message of the web interface, the built-in syslog-ng server, and other internal components to the **local** logspace.

- The second log path collects messages sent to SSB using the default syslog sources (for details, see *Section 7.1, Default message sources in SSB (p. 126)*) or via SNMP (for details, see *Procedure 7.2, Receiving SNMP messages (p. 127)*). These messages are stored in the **center** logspace.

Note
Note that both default log paths are marked as **Final**: if you create a new log path that collects logs from the default sources, make sure to adjust the order of the log paths, or disable the **Final** option for the default log path.

10.2. Procedure – Creating new log paths

Purpose:

To create a new log path, complete the following steps.

Steps:

Step 1. Navigate to **Log > Paths** and select ⊕. A new log path is added to the list of log paths.

Step 2. Select a source for the log path from the **Source** field. Messages arriving to this source will be processed by this log path. To add more sources to the log path, select ⊕ in the source field and repeat this step.

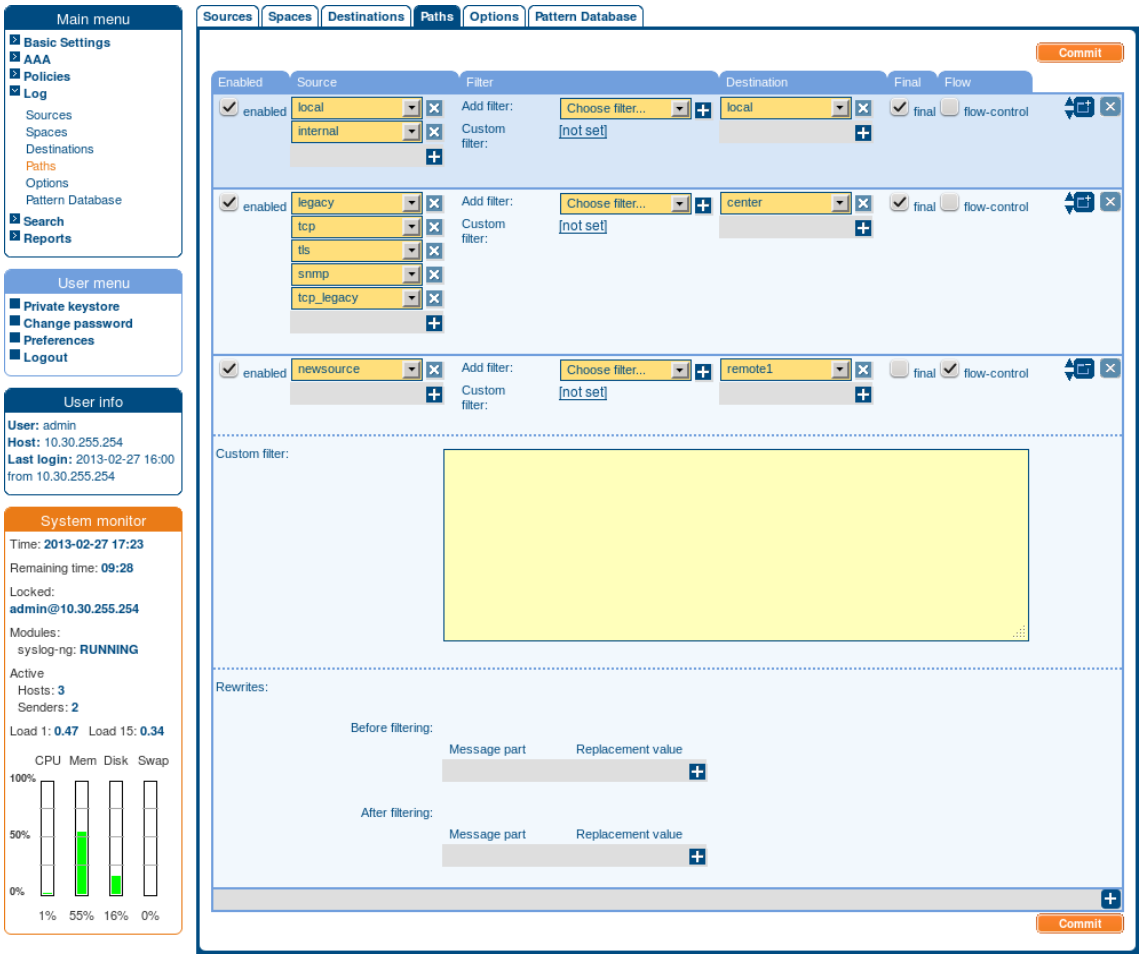

Figure 10.2. Creating a new logpath

Remote sources receive messages from the network, while *built-in* sources are messages that originate on SSB. However, note that the SNMP source (for details, see *Procedure 7.2, Receiving SNMP messages (p. 127)*) is listed in the built-in section.

Tip
To process every message of every source, leave the source option on `all`. This is equivalent to using the `catchall` flag of syslog-ng.

Step 3. Select a destination for the log path from the **Destination** field. Messages arriving to this source will be forwarded to this destination. To add more destinations to the log path, select ⊞ in the destination field and repeat this step.

Note
Remote destinations forward the messages to external servers or databases and are configured on the **Log** > **Destinations** page (for details, see *Chapter 9, Forwarding messages from SSB (p. 157)*).

Local destinations store the messages locally on SSB and are configured on the **Log** > **Spaces** page (for details, see *Chapter 8, Storing messages on SSB (p. 140)*).

If you do not want to store the messages arriving to this log path, leave the **Destination** field on *none*.

Warning
The *none* destination discards messages — messages sent only to this destination will be lost irrevocably.

Step 4. If you do not want other log paths to process the messages sent to a destination by this log path, select the **Final** option.

The order of the log paths is important, especially if you use the **Final** option in one or more destinations, because SSB evaluates log paths in descending order. Use the ⌄ buttons to position the log path if needed.

Step 5. To enable flow-control for this log path, select the **Flow** option. For details on how flow-control works, see *Section 2.3, Managing incoming and outgoing messages with flow-control (p. 5)*.

Step 6. If you do not wat to send every message from the sources to the destinations, use filters. Select the filter to use from the **Filter** field, click ⊞, and configure the filter as needed. To apply more filters, click ⊞ and select a new filter. Note that SSB sends only those messages to the destinations that pass every listed filter of the log path. The available filters are described in *Section 10.3, Filtering messages (p. 169)*.

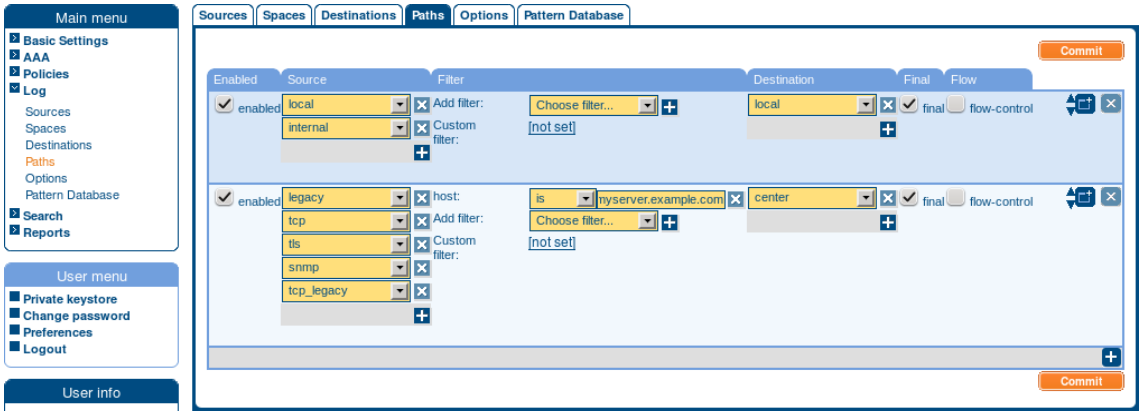

Figure 10.3. Filtering log messages

Step 7. Click Commit. After that, the new log path will start to collect log messages.

Tip
If you do not want to activate the log path immediately, deselect the **Enable** option.

10.3. Filtering messages

This section describes the filters that can be used in log paths. Every filter can be used to select (for example `priority is`) or exclude (for example `priority is not`) messages. The following filters are available:

- *facility*: Select messages sent by a specific facility (for example `kernel`).
- *host*: Select messages sent by a specific host. Enter the a hostname, IP address, or a POSIX (basic) regular expression.
- *message*: Select messages containing a specific keyword or POSIX (basic) regular expression in the text of the log message (excluding the headers).
- *priority*: Select messages of a specific priority.
- *program*: Select messages sent by a specific application. Enter the name of the application or a POSIX (basic) regular expression.
- *sender*: Filter on the address of the host that sent the message to SSB.

Note
The effect of the sender and the host filters is the same if every client sends the logs directly to SSB. But if SSB receives messages from relays, then the host filter applies for the address of the clients, while the sender applies for the address of the relays.

If multiple filters are set for a log path, only messages complying to every filter are sent to the destinations. (In other words, filters are added using the logical AND operation.)

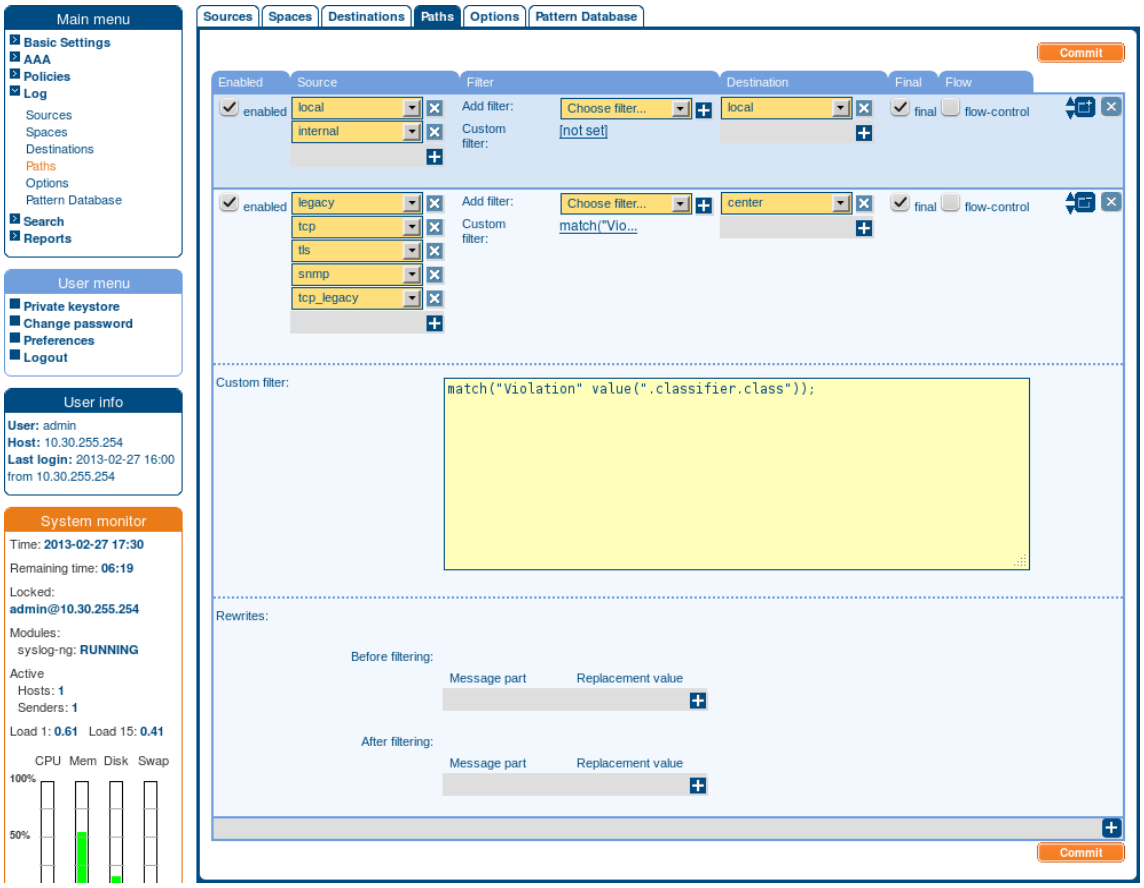

Figure 10.4. Using custom filters

If you need more complex filtering in your log path, select the ◪ of the log path and enter a custom filter into the appearing field. The contents of this field are pasted into the `filter()` parameter of the syslog-ng log path definition.

10.3.1. Procedure – Modifying messages using rewrite

Purpose:

The syslog-ng application can rewrite parts of the messages using rewrite rules. Almost all parts of the message can be rewritten. The rules use a key-value pair format.

Steps:

Step 1. Navigate to **Log** > **Paths**.

Step 2. Select the path(s) where you want to use rewrite rules.

Step 3. In the **Rewrites** section, click ⊞ to add a new rewrite rule. Rewrite rules can be applied before filtering, or after filtering.

The sequence of filtering and rewrite rules depends on how it was specified in the log path. The sequence of the process is the following:

1. Rewrite the message parts using the "before filtering" rewrite rules in the order the rewrite rules were given.
2. Filter the messages.
3. Rewrite the message parts using the "after filtering" rewrite rules in the order the rewrite rules were given.
4. Send the messages to the given destinations.

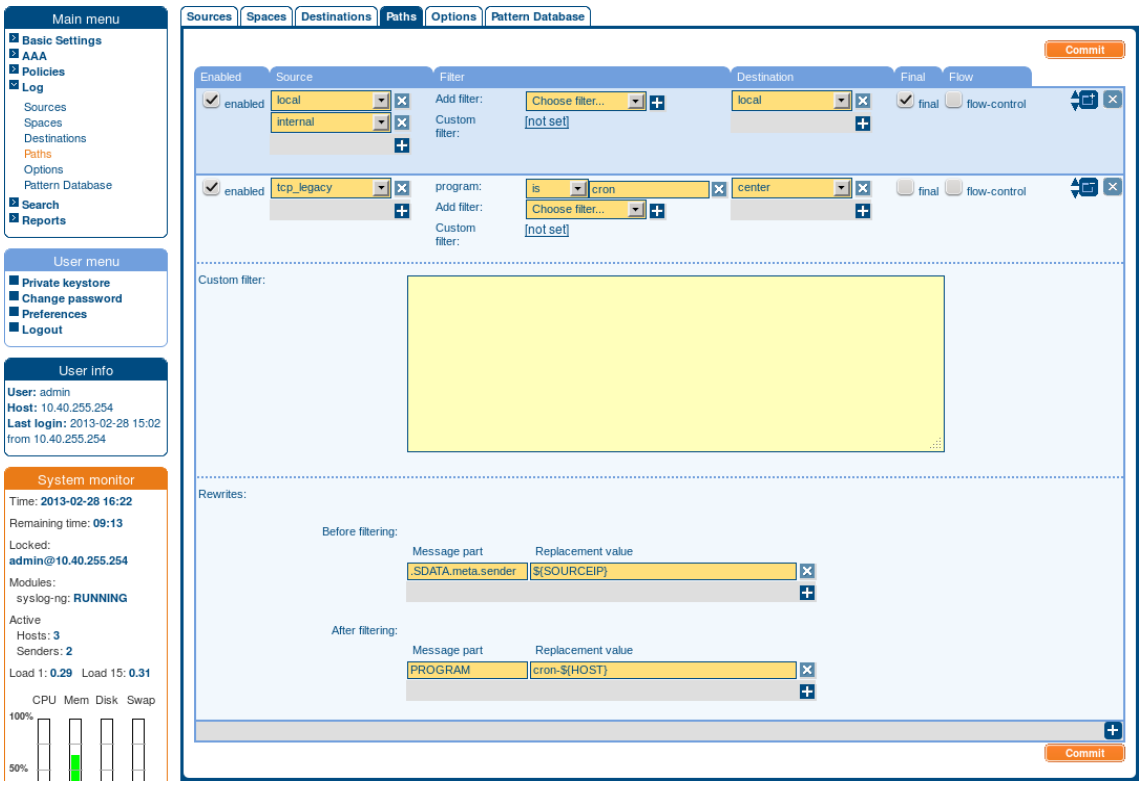

Figure 10.5. Modifying messages using rewrite

Step 4. Enter the part of the message to rewrite in the **Message part** field. For example *`MESSAGE`*; *`HOST`*; *`.SDATA.meta.custom`*. If the specified field does not exist, it is automatically created and set to the **Replacement value** field.

Step 5. Enter the value of the message part after rewriting in the **Replacement value** field. To use macros, begin with a $ sign and enclose the name of the macro between braces, for example *`${MSG}`*; *`${.SDATA.meta.custom}`*.

Note
The replacement value completely replaces the old value of the message part.

Note
Hard macros contain data that is directly derived from the log message. It is not possible to change the values of hard macros in rewrite rules. For the list of hard macros, see *Section Hard vs. soft macros* in *The syslog-ng Premium Edition 5 LTS Administrator Guide*.

Chapter 11. Configuring syslog-ng options

There are several options of the syslog-ng server running on SSB that can be configured. These include:

- For details on general syslog-ng settings — see *Section 11.1, General syslog-ng settings (p. 173)*.
- For details on timestamping-related options — see *Section 11.2, Timestamping configuration on SSB (p. 175)*.
- For details on certificate management for receiving and sending log messages in TLS-encrypted channels — see *Procedure 11.4, Setting the certificates used in TLS-encrypted log transport (p. 177)*.
- For details on managing domain name resolution for log messages — see *Section 11.3, Using name resolution on SSB (p. 176)*.

11.1. General syslog-ng settings

To configure the general options of the syslog-ng server running on SSB, navigate to **Log > Options**. The following options are available (note that options related to name resolution are discussed in *Section 11.3, Using name resolution on SSB (p. 176)*):

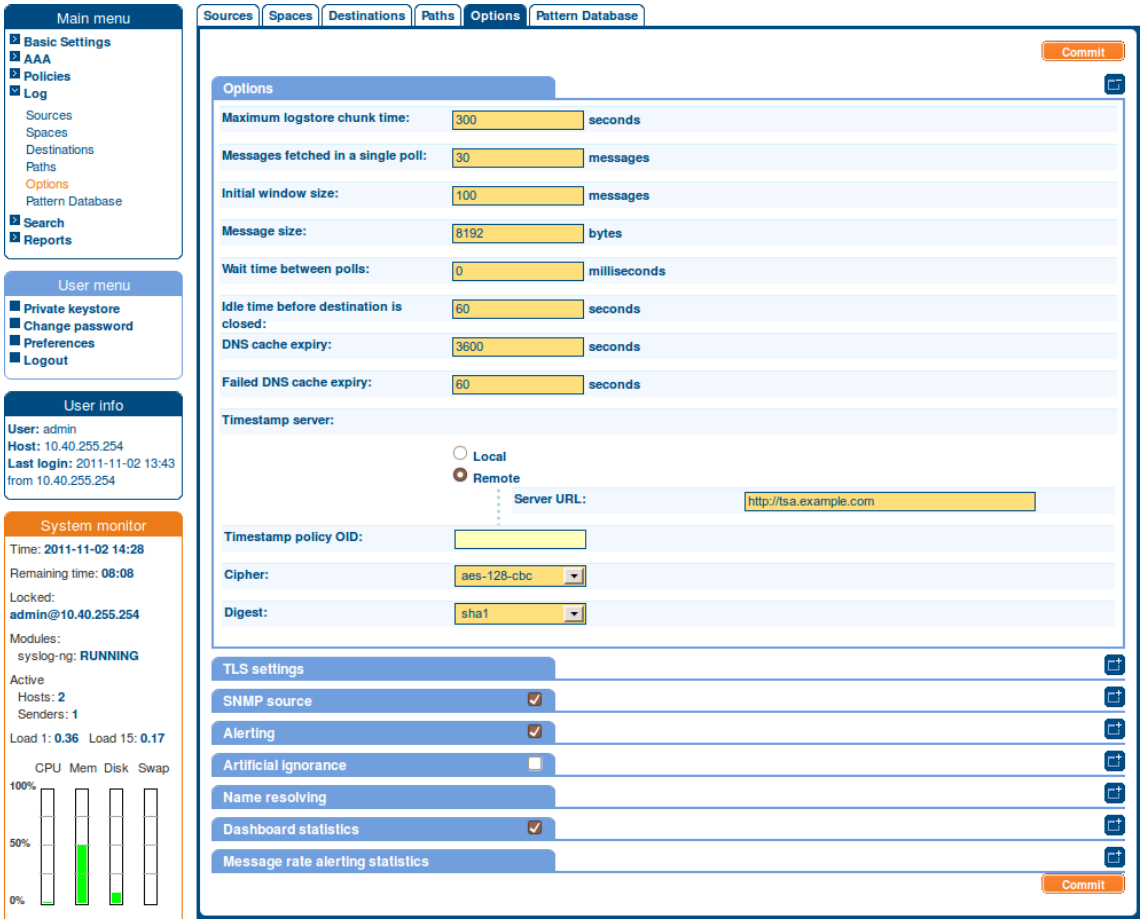

Figure 11.1. Configuring syslog-ng options

- *Maximum logstore chunk time*: Time limit in seconds: syslog-ng closes the chunk if no new messages arrive until the time limit expires. Logstore chunks are closed when the time limit expires. If the time limit set in the **Idle time before destination is closed** option expires, the entire file is closed. This option corresponds to the `chunk_time()` parameter of syslog-ng.

- *Messages fetched in a single poll*: The maximum number of messages fetched from a source during a single poll loop. The destination queues might fill up before flow-control could stop reading if this parameter is too high. This option corresponds to the `log_fetch_limit()` parameter of syslog-ng.
- *Initial window size*: The size of the initial window used during flow control. This option corresponds to the `log_iw_size()` parameter of syslog-ng.
- *Message size*: Specifies the maximum length of incoming log messages in bytes. This option corresponds to the `log_msg_size()` parameter of syslog-ng. The maximum value of this parameter is 1000000 (1 MB).
- *Wait time between polls*: The time to wait in milliseconds before checking if new messages have arrived to a source. This option corresponds to the `time_sleep()` parameter of syslog-ng.
- *Idle time before destination is closed*: The time to wait in seconds before an idle destination file is closed. This option corresponds to the `time_reap()` parameter of syslog-ng.

11.2. Timestamping configuration on SSB

To configure the timestamping options of SSB, navigate to **Log > Options**. The following options are available:

- *Timestamp server*: Select the timestamping server to use for signing encrypted logspaces. To use the built-in timestamp server of SSB, select **Local**.

  To use an external timestamping server, select **Remote** and enter the address of the server into the **Server URL** field. Note that currently only plain HTTP services are supported, password-protected and HTTPS services are not supported at.

Warning
SSB currently supports only timestamping servers that use the *Internet X.509 Public Key Infrastructure Time-Stamp Protocol (TSP)* described in RFC 3161.

- *Timestamp policy OID*: If the Timestamping Server has timestamping policies configured, enter the OID of the policy to use into the Timestamping policy field. SSB will include this ID in the timestamping requests sent to the TSA.

- *Cipher*: Select the cipher method used to encrypt the logstore. The following cipher methods are available: `aes-128-cbc`, `aes-128-cfb`, `aes-128-cfb1`, `aes-128-cfb8`, `aes-128-ecb`, `aes-128-ofb` , `aes-192-cbc`, `aes-192-cfb`, `aes-192-cfb1`, `aes-192-cfb8`, `aes-192-ecb`, `aes-192-ofb` , `aes-256-cbc`, `aes-256-cfb`, `aes-256-cfb1`, `aes-256-cfb8`, `aes-256-ecb`, `aes-256-ofb` , `aes128` , `aes192` , `aes256`, `bf` , `bf-cbc` , `bf-cfb`, `bf-ecb` , `bf-ofb` , `blowfish`, `cast` , `cast-cbc` , `cast5-cbc` , `cast5-cfb`, `cast5-ecb`, `cast5-ofb` , `des`, `des-cbc`, `des-cfb` , `des-cfb1` , `des-cfb8` , `des-ecb` , `des-ede`, `des-ede-cbc`, `des-ede-cfb` , `des-ede-ofb`, `des-ede3` , `des-ede3-cbc`, `des-ede3-cfb`, `des-ede3-ofb`, `des-ofb` , `des3` , `desx` , `desx-cbc`, `rc2`, `rc2-40-cbc` , `rc2-64-cbc`, `rc2-cbc`, `rc2-cfb`, `rc2-ecb` , `rc2-ofb`, `rc4`, and `rc4-40`.

  By default, SSB uses the `aes-128-cbc` method.

- *Digest*: Select the digest method to use. The following digest methods are available: `MD2`, `MD4`, `MD5`, `SHA-0 (SHA)`, `SHA-1`, `RIPEMD-160`, `SHA-224`, `SHA-256`, `SHA-384`, and `SHA-512`.

  By default, SSB uses the `SHA-1` method.

Warning
The size of the digest hash must be equal to or larger than the key size of the cipher method. For example, to use the `aes-256-cbc` cipher method, the digest method must be at least `SHA-256`.

Note
The timestamp requests are handled by a separate process in syslog-ng; message processing is not affected if the timestamping server is slow or cannot be accessed.

11.3. Using name resolution on SSB

SSB can resolve the hostnames of the clients and include them in the log messages. However, the performance of SSB can be severely degraded if the domain name server is unaccessible or slow. Therefore, SSB automatically caches the results of name resolution. If you experience performance problems under high load, it is not recommended to disable name resolution. If you must use name resolution, consider the following:

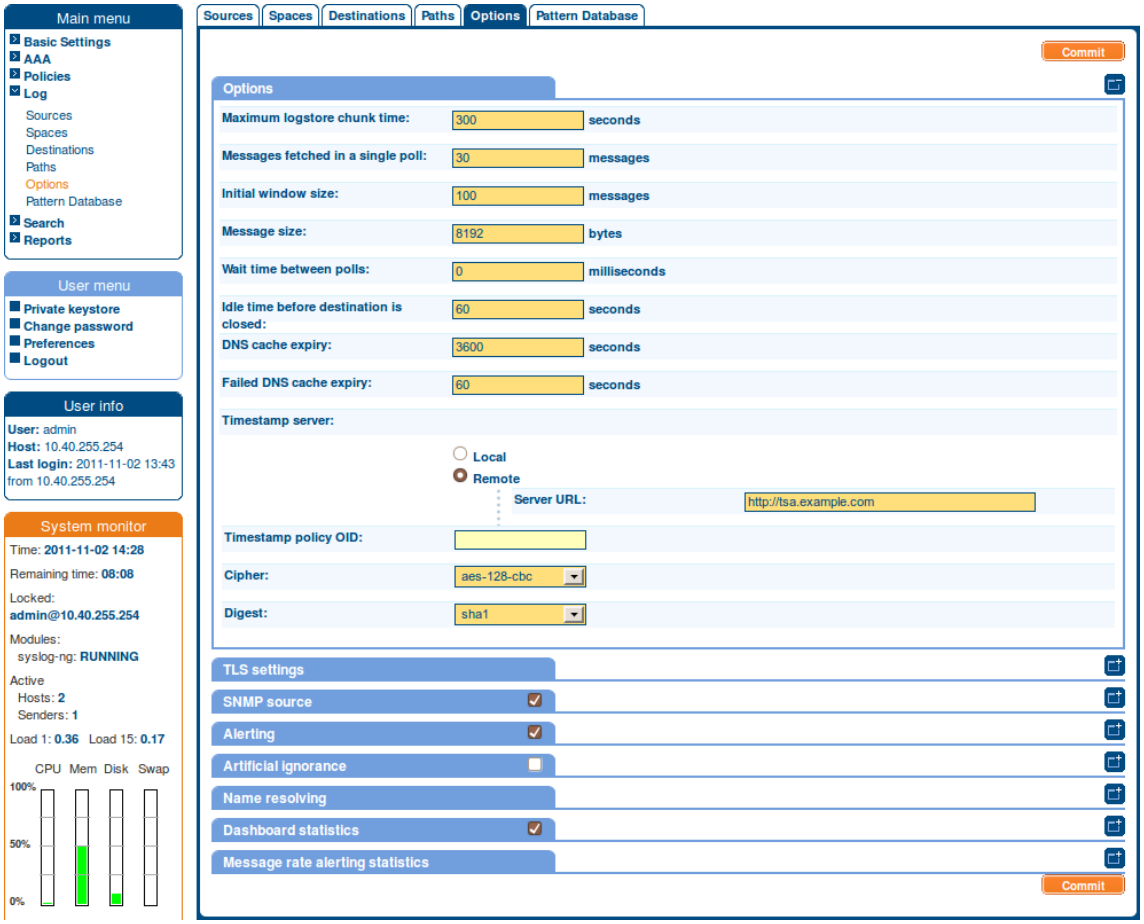

Figure 11.2. Configuring DNS options

- If the IP addresses of the clients change only rarely, set the expiry of the DNS cache to a large value. By default, SSB caches successful DNS lookups for an hour, and failed lookups for one minute. These parameters can be adjusted under **Log > Options > Options > DNS Cache expiry** and **Failed DNS cache expiry**.

- Resolve the hostnames locally. Resolving hostnames locally enables you to display hostnames in the log files for frequently used hosts, without having to rely on a DNS server. The known IP address – hostname pairs are stored locally in a file. In the log messages, syslog-ng will replace the IP

addresses of known hosts with their hostnames. To configure local name resolution, select **Log** > **Options** > **Name resolving**, and enter the IP Address - hostname pairs in (for example *`192.168.1.1 myhost.example.com`*) into the **Persistent hostname list** field. Then navigate to **Log** > **Sources**, and set the **Use DNS** option of your sources to **Only from persistent configuration**.

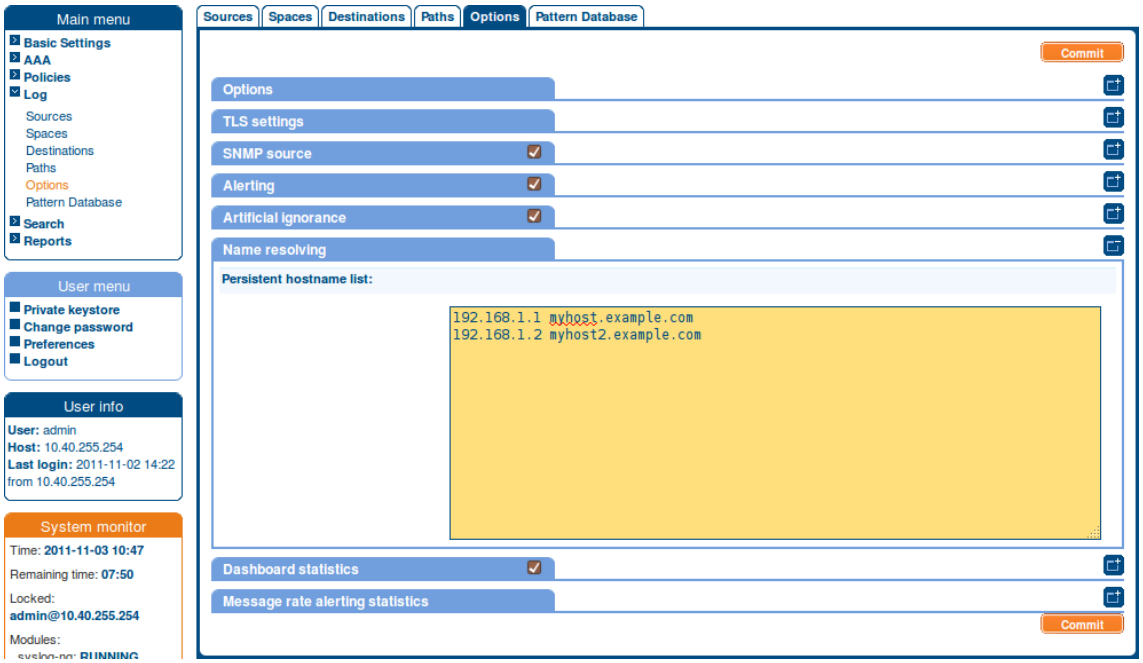

Figure 11.3. Configuring persistent name resolution

11.4. Procedure – Setting the certificates used in TLS-encrypted log transport

Purpose:

To set a custom certificate and a CA certificate for encrypting the transfer of log messages, complete the following steps.

Note
If you do not upload a certificate to encrypt the TLS-communication (that is, the **TLS certificate** and **TLS private key** options are not set), SSB uses the certificate and CA certificate set for the web interface (set under **Basic Settings** > **Management** > **SSL certificates**) for this purpose as well.

BalaBit recommends using 2048-bit RSA keys (or stronger).

Steps:

Step 1. In your PKI system, generate and sign a certificate for SSB, then navigate to **Log** > **Options** > **TLS settings**.

Step 2. Click the icon in the **TLS certificate** field to upload the certificate.

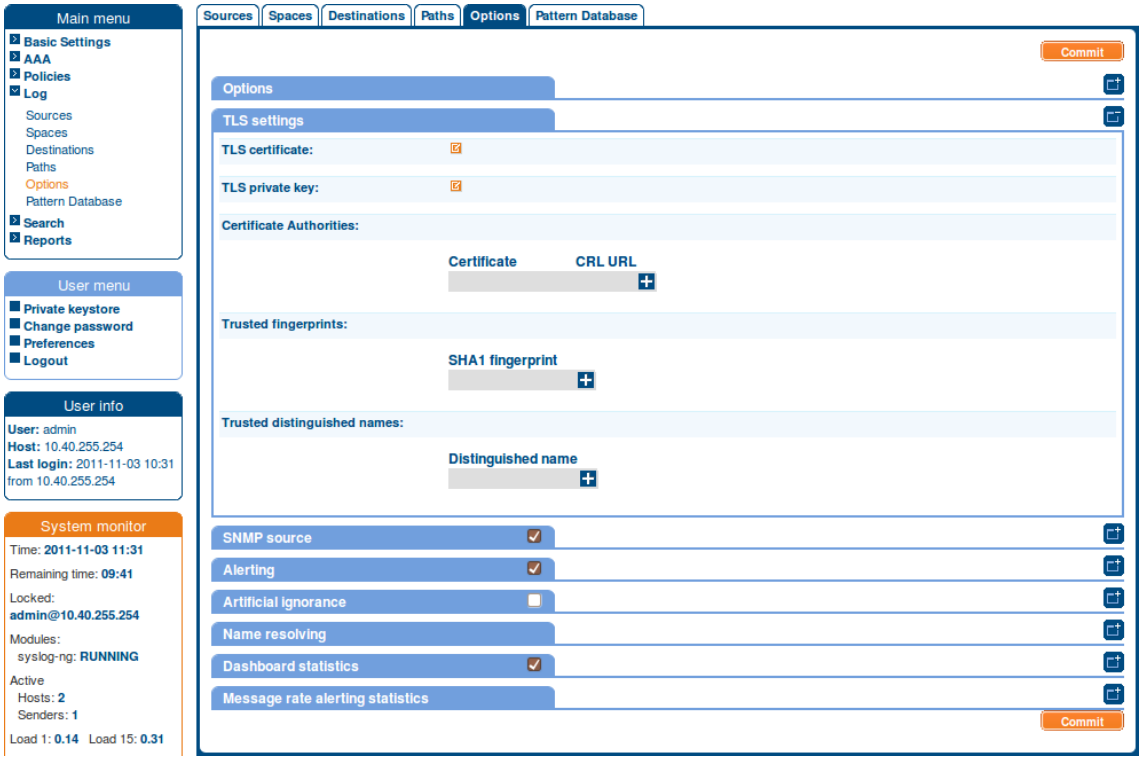

Figure 11.4. Configuring TLS settings for syslog-ng

To upload a certificate from a file, click **Browse** in the **Upload key** section, select the certificate file, and click **Upload**.

Alternatively, you can copy/paste the certificate into the **Key** field of the **Copy-paste key** section and click **Upload**.

Step 3. Click the icon in the **TLS private key** field and upload the private key corresponding to the certificate.

Step 4. To set the certificate of the Certificate Authority (CA) used to verify the identity of the peers, click in the **Certificate Authorities** field, then click .

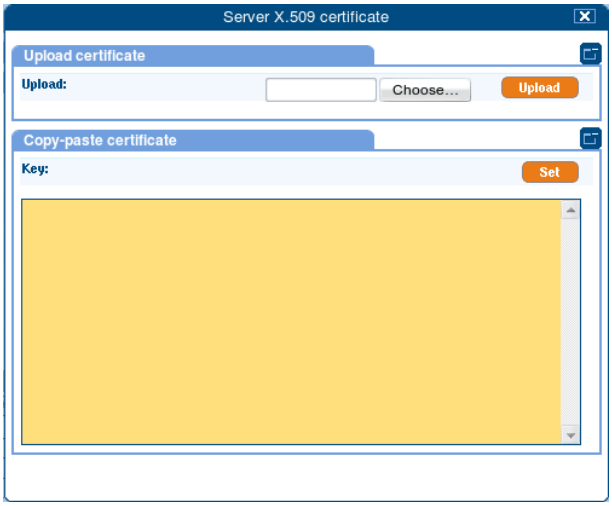

Figure 11.5. Uploading certificates

To upload a certificate from a file, click **Browse** in the **Upload key** section, select the certificate file, and click **Upload**.

Alternatively, you can copy/paste the certificate into the **Key** field of the **Copy-paste key** section and click **Upload**.

Repeat this step to add more CA certificates if needed.

Step 5. If the CA issues a Certificate Revocation List (CRL), enter its URL into the **CRL URL** field. SSB periodically downloads the list and refuses certificates that appear on the list.

Note
Note that only *.pem* format CRLs are accepted. CRLs that are in PKCS7 format (*.crl*) are not accepted.

Step 6. If you want to accept connections only from hosts using certain certificates signed by the CA, click ⊞ in the **Trusted distinguished names** field and enter the distinguished name (DN) of the accepted certificates into the **Distinguished name** field. This option corresponds to the *trusted-dn()* parameter of syslog-ng.

Example: `*, O=Example Inc, ST=Some-State, C=*` accepts only certificates issued for the *Example Inc* organization in *Some-State* state.

Step 7. If you want to accept a certificate without uploading its corresponding CA certificate, click ⊞ in the **Trusted fingerprints** field and enter the SHA-1 fingerprint of the accepted certificates into the **SHA-1 fingerprint** field. This option corresponds to the *trusted-keys()* parameter of syslog-ng.

Example: `SHA1:00:EF:ED:A4:CE:00:D1:14:A4:AB:43:00:EF:00:91:85:FF:89:28:8F, SHA1:0C:42:00:3E:B2:60:36:64:00:E2:83:F0:80:46:AD:00:A8:9D:00:15` adds these specific SHA-1 fingerprints:

SHA1:00:EF:ED:A4:CE:00:D1:14:A4:AB:43:00:EF:00:91:85:FF:89:28:8F and *SHA1:0C:42:00:3E:B2:60:36:64:00:E2:83:F0:80:46:AD:00:A8:9D:00:15*.

Chapter 12. Browsing log messages

This section describes how to browse the log messages collected on SSB.

- *Section 12.1, Using the search interface (p. 181)* explains how to use and customize the search interface, describes the log message data that is available on SSB, and provides examples of the the wildcard and boolean search operators you can use.
- *Section 12.2, Browsing encrypted log spaces (p. 190)* describes how to decrypt and browse encrypted logspaces.
- *Section 12.3, Creating custom statistics from log data (p. 195)* explains how to create custom statistics from the available log data, and how to save them for reports.

12.1. Using the search interface

SSB has a search interface for browsing the collected log messages. You can choose the logspace, enter a search query, specify the timeframe, and browse the results here.

To access the search interface, navigate to **Search > Spaces**.

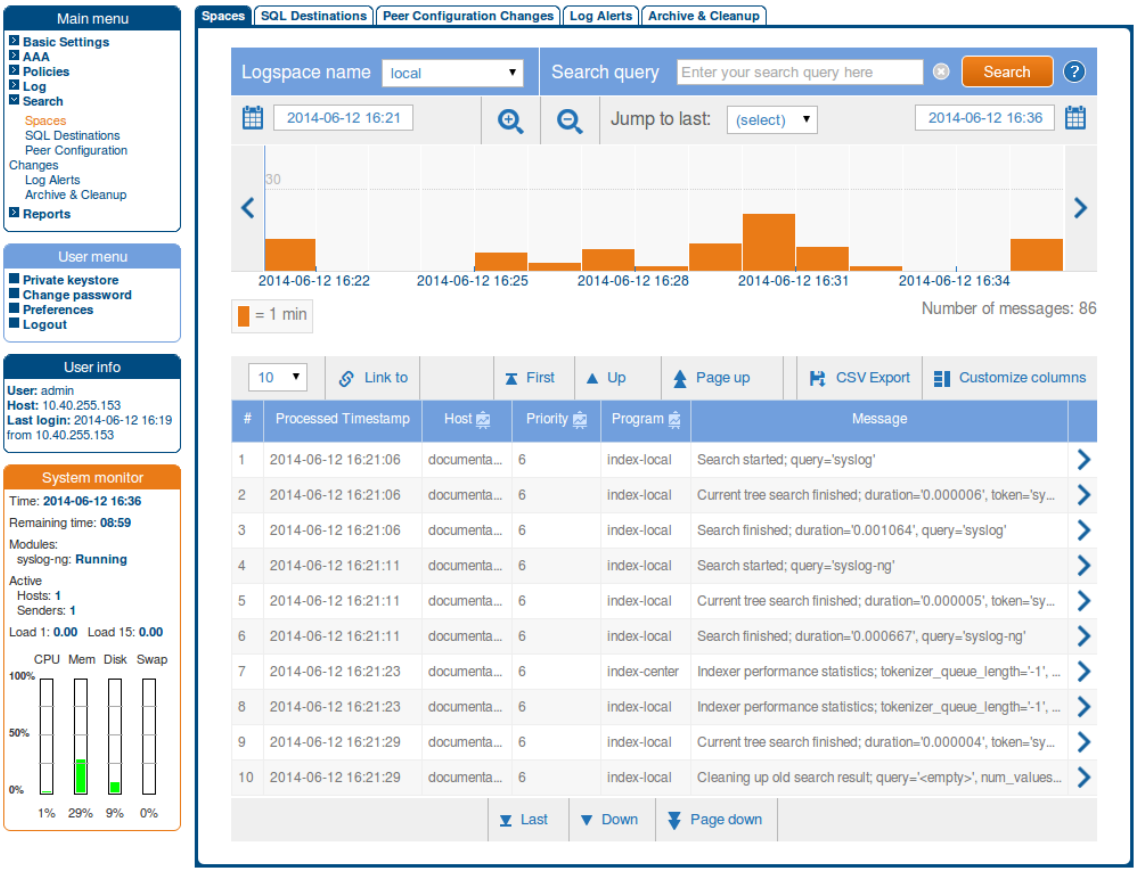

Figure 12.1. The log message search interface

Logspaces:

Choose the appropriate logspace using the **Logspace name** menu.

For more information on the available logspaces, and how to configure them, see *Chapter 8, Storing messages on SSB (p. 140)*.

Search:

On the log message search interface, you can use the Search Query field to search the full list of log messages. Click on the icon, or see *Section 12.1.3, Using wildcards and boolean search (p. 187)* for more details.

When searching log messages, the capabilities of the search engine depend on the delimiters used to index the particular log space. For details on how to configure the delimiters used for indexing, see *Procedure 8.4.1, Creating a new logstore (p. 143)*.

You can create complex searches using wildcards and boolean expressions. For more information and practical examples, see *Section 12.1.3, Using wildcards and boolean search (p. 187)*.

Note
Note that SSB only indexes the first 60 characters of every name of a name-value pair. If the name is longer than 60 characters, an exact search does not give any results.

For example, if the name of a name-value pair is `.sdata.security.uid=2011-12-08T12:32:25.024+01:00-hostname-12345`, SSB will index only the first 60 characters: `.sdata.security.uid=2011-12-08T12:32:25.024+01:00-hostname-`. Searching for `.sdata.security.uid=2011-12-08T12:32:25.024+01:00-hostname-12345`, returns no results, so search for `.sdata.security.uid=2011-12-08T12:32:25.024+01:00-hostname-` (or a part of it) instead.

Overview:

Displays the number of log messages in the selected time interval.

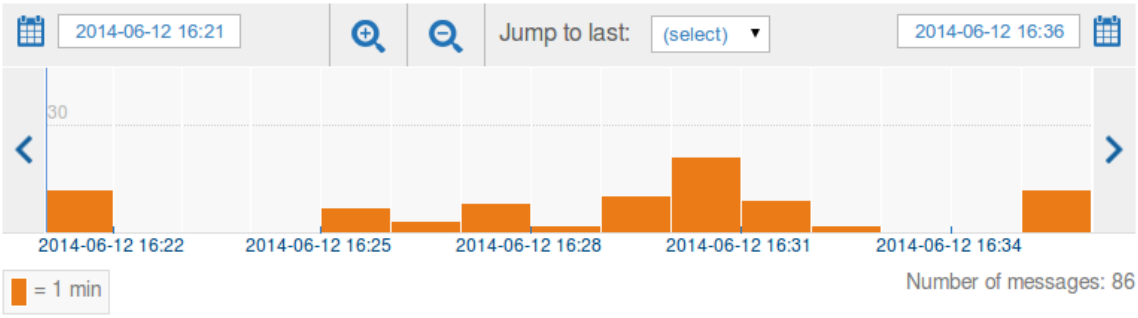

Figure 12.2. Log message overview

Use the , icons to zoom, and the arrows to display the previous or the next intervals. To change the timeframe, you can:

- Change the beginning and end date.
- Click and drag the pointer across a period on the calendar bars to select a specific interval and zoom in.
- Use the **Jump to last option** to select the last 15 minutes, hour, 6 hours, day, or week.

Hovering the mouse above a bar displays the number of entries, and the start and end date of the period that the bar represents. Click a bar to display the entries of that period in the table. Use Shift+Click to select multiple bars.

List of log messages:

Use the arrow keys and Page Up and Page Down to navigate the listed log messages, or enable mouse wheel scrolling in your Preferences. If data is too long to fit on one line, it is automatically wrapped and only the first line is displayed.

- To expand a row in the log message search interface, click ›. The complete log message is displayed:

Message 8 of 86
Processed timestamp: 2014-06-12 16:21:23
Timestamp: 2014-06-12 16:21:23
Host: documentation-ssb
Program[PID]: index-local[3375]
Facility: 5
Priority: 6
Unique ID: 808576453019959431
Tags:
Message: Indexer performance statistics; tokenizer_queue_length='-1', tokenizer_queue_histogram='0,0,0,0,0', number_of_receiver_waits='0', number_of_processed_messages='130', size_of_processed_messages='104252', average_message_size='801', average_tokens_per_message='10', number_of_receiver_reads='126', receiver_last_message_timestamp='1402582871', receiver_last_message_delay_to_wall_clock='12', tokenizer_last_message_timestamp='1402582871', tokenizer_last_message_delay_to_wall_clock='12', tokenizer_queue_delay='0', current_tree_number_of_msg_ids='1188', current_tree_number_of_tokens='211', current_tree_number_of_nodes='493', current_tree_memory_usage='42520'
Dynamic columns: .sdata.timequality.issynced=0

Figure 12.3. Viewing a single log message

To return to the list of all log messages, click <.

12.1.1. Procedure – Customizing columns of the log message search interface

To customize the data displayed on the log message search interface, complete the following steps:

Steps:

Step 1. Click **Customize columns**.

Step 2. The displayed parameters are enlisted in the **Displayed columns** field. All other available parameters are enlisted in the **Available static columns** and **Available dynamic columns** fields.

Dynamic columns are created from name-value pairs.

Note
To export the search results into a CSV file, click **Search > CSV Export**. Note that the CSV file includes all the static columns and the displayed dynamic columns.

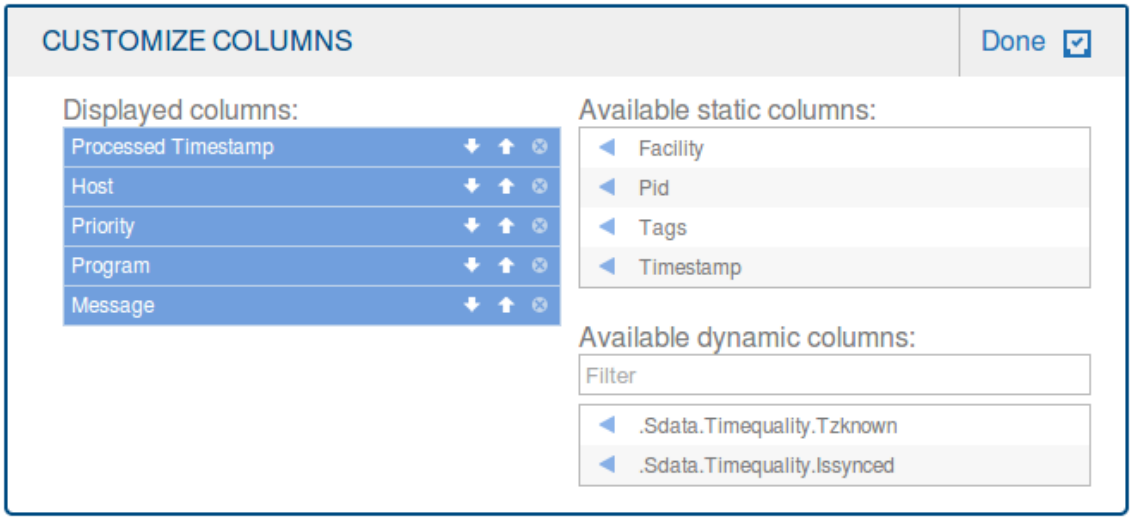

Figure 12.4. Customizing columns of the log message search interface

- To add a static column to the **Displayed columns**, click ◂.
- To add a dynamic column to the **Displayed columns**, choose a name-value pair from **Available dynamic columns** and click ◂.

  The selected name generates a new, separate dynamic column with **<name>** heading (where **<name>** is the name of the key). The relevant values are displayed in the cells of the respective column.
- To remove parameters from the **Visible columns**, click ⊠.

12.1.2. Metadata collected about log messages

The following information is available about the log messages:

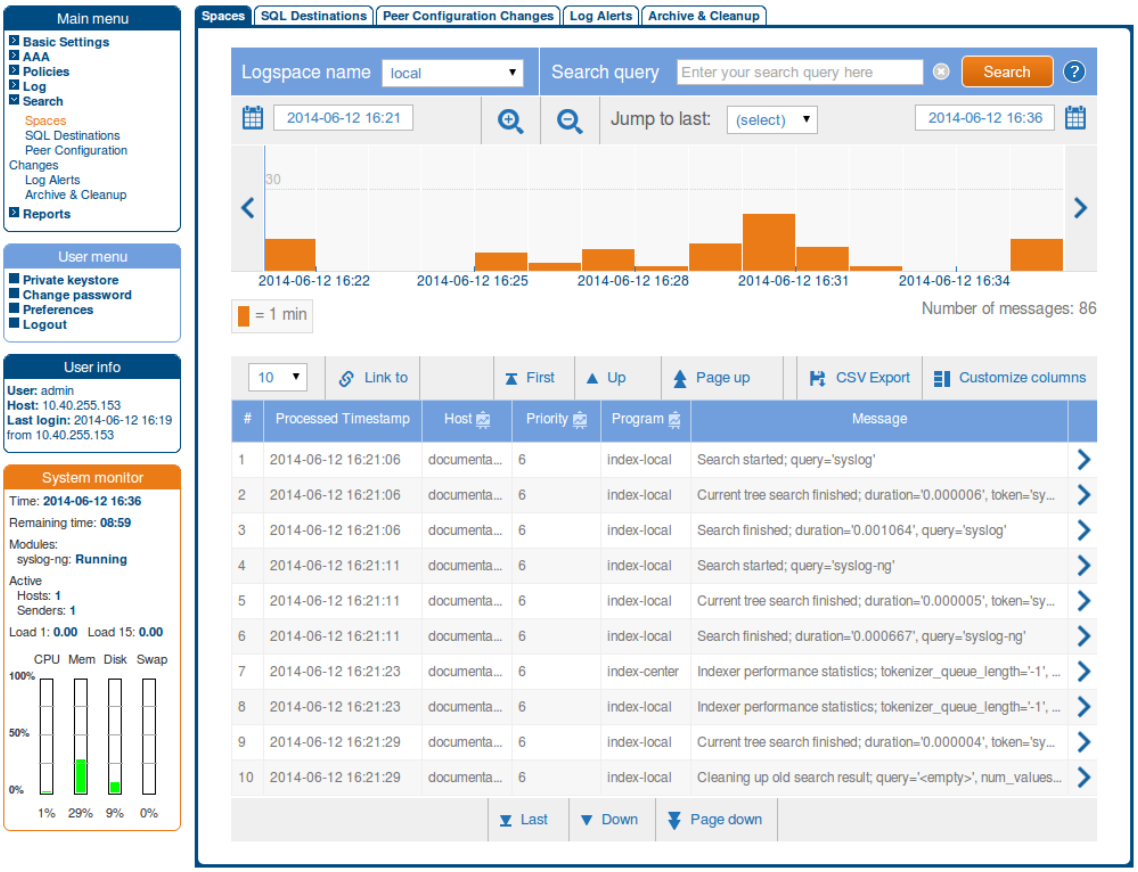

Figure 12.5. Displaying search information

- *Processed Timestamp*: The date when SSB has received the log message in `YEAR-MONTH-DAY HOUR:MINUTE:SECOND` format.
- *Timestamp*: The timestamp received in the message — the time when the log message was created in `YEAR-MONTH-DAY HOUR:MINUTE:SECOND` format.
- *Facility*: The facility that sent the message.
- *Priority*: The priority value of the message.
- *Program*: The application that created the message.
- *Pid*: The program identifier of the application that created the message.
- *Host*: The IP address or hostname of the client that sent the message to SSB.
- *Message*: The text of the log message.
- *Tag*: Tags assigned to the message matching certain pattern database rules.
- *Id*: Unique ID of the message.
- *classifier.rule_id*: ID of the pattern database rule that matched the message.
- *classifier.class*: Description of the pattern database rule that matched the message.
- Dynamic columns, created from additional name-value pairs, might also be available.

12.1.3. Using wildcards and boolean search

You can use wildcards and boolean expressions to search the log messages collected on SSB. The following sections provide examples for different search queries.

Note
Note that SSB only indexes the first 60 characters of every name of a name-value pair. If the name is longer than 60 characters, an exact search does not give any results.

For example, if the name of a name-value pair is `.sdata.security.uid=2011-12-08T12:32:25.024+01:00-hostname-12345`, SSB will index only the first 60 characters: `.sdata.security.uid=2011-12-08T12:32:25.024+01:00-hostname-`. Searching for `.sdata.security.uid=2011-12-08T12:32:25.024+01:00-hostname-12345`, returns no results, so search for `.sdata.security.uid=2011-12-08T12:32:25.024+01:00-hostname-` (or a part of it) instead.

- For examples of exact matches, see *Section Searching for exact matches and using complex queries (p. 187)*.
- For examples of searching in a specific part of the message, see *Section Searching in a specific part of the message (p. 187)*.
- For examples of using boolean operators to combine search keywords, see *Section Combining search keywords (p. 188)*.
- For examples of wildcard searches, see *Section Using wildcard searches (p. 188)*.
- For examples of searching with special characters, see *Section Searching for special characters (p. 190)*.

Searching for exact matches and using complex queries

By default, SSB searches for keywords as whole words in the `MESSAGE` part of the log message and returns only exact matches.

Example 12.1. Searching for exact matches

Search expression	`example`
Matches	example
Does not match	examples
	example.com
	query-by-example
	exam

Searching in a specific part of the message

You can search in a specific part of the message using the `<type>:` prefix. The `message:` (or `msg:`) prefix means the message part and can be omitted. For example: to search for the name of an application, use the `program:` prefix; to search for a host name, use the `host:` prefix, and so on.

Example 12.2. Searching specific parts of messages

Search expression	`program:syslog-ng`
Matches	All log messages from the syslog-ng application.

Combining search keywords

You can use boolean operators – *AND*, *OR*, and *NOT* – to combine search keywords. Note that the boolean operators are case sensitive, and must be in all caps. More complex search expressions can also be constructed with parentheses.

Example 12.3. Combining keywords in search

Search expression	`keyword1 AND keyword2`
Matches	(returns log messages that contain both keywords)
Search expression	`keyword1 OR keyword2`
Matches	(returns log messages that contain at least one of the keywords)
Search expression	`keyword1 AND NOT keyword2`
Matches	(returns log messages that contain only `keyword1`)

To search for expressions that can be interpreted as boolean operators (for example: *AND*), use the following format: *message:AND*.

Example 12.4. Using parentheses in search
Use parentheses to create more complex search expressions:

Search expression	`(keyword1 OR keyword2) AND keyword3`
Matches	(returns log messages that contain either keyword1 and keyword3, or keyword2 and keyword3)

Using wildcard searches

You can use the *?* and *** wildcards in your search expressions.

Example 12.5. Using wildcard ? in search
The *?* (question mark) wildcard means exactly one arbitrary character. Note that it does not work for finding non-UTF-8 or multibyte characters. If you want to search for these characters, the expression *??* might work, or you can use the *** wildcard instead.

Search expression	`example?`
Matches	example1
	examples
Does not match	example.com
	example12
	query-by-example
	example?

Search expression	`?example?`
Matches	1example2
Does not match	example.com
	example12
	query-by-example
Search expression	`example??`
Matches	example12
Does not match	example.com
	example1
	query-by-example

Example 12.6. Using wildcard * in search
The `*` wildcard means 0 or more arbitrary characters. It finds non-UTF-8 and multibyte characters as well. Wildcard characters also work in any message part, for example, *`program:postfix*`*.

Search expression	`example*`
Matches	example
	examples
	example.com
Does not match	query-by-example
	example*
Search expression	`*example`
Matches	example
	query-by-example
	example.com
Does not match	example.com
	example12
Search expression	`*example*`
Matches	example
	query-by-example
	example.com
	example12
Does not match	example.com
	example12

Example 12.7. Using combined wildcards in search
Wildcard characters can be combined.

Search expression	`ex?mple*`
Matches	example1
	examples
	example.com
	exemple.com
Does not match	exmples
	example12
	query-by-example

Searching for special characters

To search for the question mark (*?*), asterisk (***), backslash (\) or whitespace ( ) characters, you must prefix these characters with a backslash (\). Any character after a backslash is handled as character to be searched for.

Example 12.8. Searching for special characters
To search for a special character, use a backslash (\).

Search expression	`example\?`
Matches	example?
Does not match	examples
	example1

To search for a special character, backslash character, use two backslashes (\\).

Search expression	`C:\\Windows`
Matches	C:\Windows

Search expression	`nvpair:path=C:\\Program\ Files`
Matches	C:\Program Files

Note
It is not possible to search for whitespace ( ) character in the MESSAGE part of the log message, since it is a hard-coded delimiter character.

12.2. Browsing encrypted log spaces

By default, you cannot browse encrypted logstores from the SSB web interface, because the required decryption keys are not available on SSB. To make browsing and searching encrypted logstores possible, SSB provides the following options:

- Use persistent decryption key(s) for a single user.
  For details, see *Procedure 12.2.1, Using persistent decryption keys (p. 191)*.

- Use decryption keys for the duration of the user session only.
  For details, see *Procedure 12.2.2, Using session-only decryption keys (p. 193)*.

- Assign decryption keys to a logstore (making them available to every SSB user). This option might raise security concerns.
  For details, see *Procedure 12.2.3, Assigning decryption keys to a logstore (p. 194)*.

Note
Do not use SSB's own keys and certificates for encrypting or decrypting.

BalaBit recommends using 2048-bit RSA keys (or stronger).

12.2.1. Procedure – Using persistent decryption keys

Purpose:

You can upload decryption keys and bind them to your account. The decryption keys are stored on SSB, but they are only made available for this user account, and can also be protected (encrypted) with a passphrase.

Steps:

Step 1. Select **User Menu** > **Private keystore**. A popup window is displayed.

Step 2. Select **Permanent** > , then select **Certificate** > . A popup window is displayed.

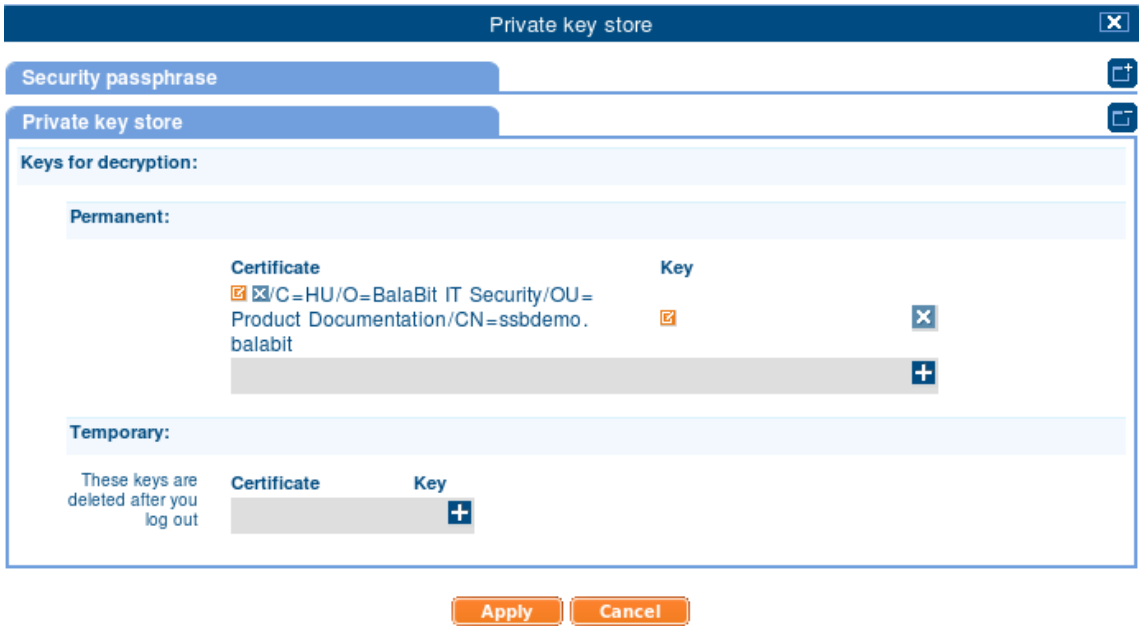

Figure 12.6. Adding decryption keys to the private keystore

Step 3. Paste or upload the certificate used to encrypt the logstore.

Step 4. Select **Key >** . A popup window is displayed.

Step 5. Paste or upload the private key of the certificate used to encrypt the logstore.

Step 6. Repeat Steps 2-5 to upload additional keys if needed.

Step 7. Select **Security passphrase > Change**, and enter a passphrase to protect the private keys.

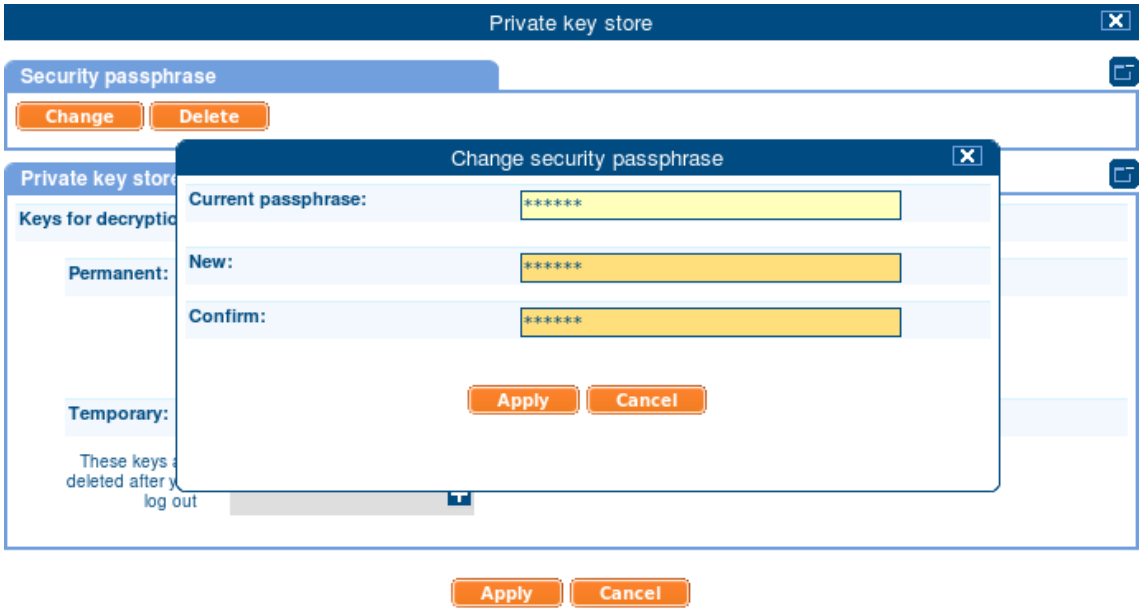

Figure 12.7. Securing the private keystore with a passphrase

Step 8. Click **Apply**.

12.2.2. Procedure – Using session-only decryption keys

Purpose:

You can upload decryption keys to browse encrypted logspaces for the duration of the session only. These keys are automatically deleted when you log out from SSB.

Steps:

Step 1. Select **User Menu** > **Private keystore**. A popup window is displayed.

Step 2. Select **Temporary** > , then select **Certificate** > . A popup window is displayed.

Figure 12.8. Adding decryption keys to the private keystore

Step 3. Paste or upload the certificate used to encrypt the logstore.

Step 4. Select **Key** > . A popup window is displayed.

Step 5. Paste or upload the private key of the certificate used to encrypt the logstore.

Step 6. Repeat Steps 2-5 to upload additional keys if needed.

Step 7. Click **Apply**.

12.2.3. Procedure – Assigning decryption keys to a logstore

Purpose:

You can add a private key (or set of keys) to a logstore, and use these keys to decrypt the logstore files. This way, anyone who has the right to search the particular logspace can search the messages. These decryption keys are stored unencrypted in the SSB configuration file.

As this may raise security concerns, avoid this solution unless absolutely necessary.

Steps:

Step 1. Navigate to **Log** > **Spaces** and select the encrypted logspace you want to make searchable for every user via the SSB web interface.

Step 2. Select **Decryption private keys** > . A popup window is displayed.

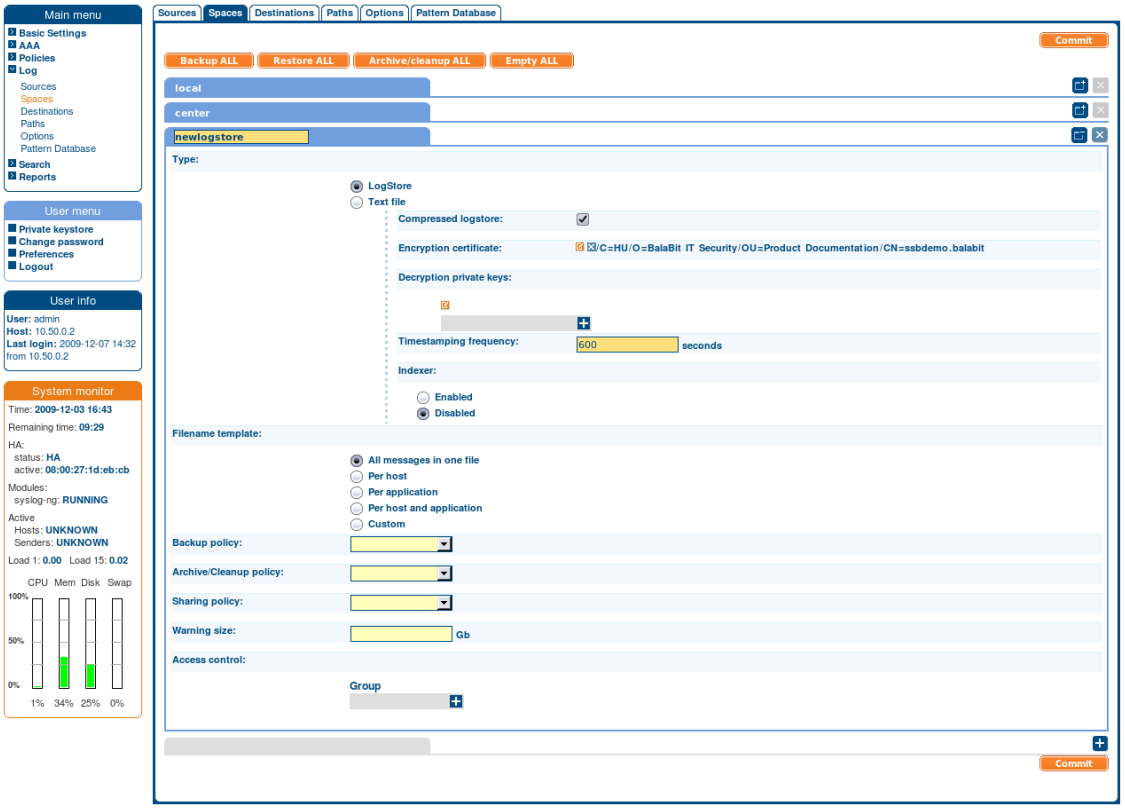

Figure 12.9. Adding decryption keys to a logstore

Step 3. Paste or upload the private key of the certificate used to encrypt the logstore.

Step 4. Repeat Steps 2-3 to upload additional keys if needed.

Step 5. Click Commit.

12.3. Creating custom statistics from log data

SSB can create statistics from the `Timestamp`, `Facility`, `Priority`, `Program`, `Pid`, `Host`, `Tags`, and `.classifier.class` columns. Use **Customize columns** to add the required column, if necessary.

To display statistics about the log messages, click the icon in the appropriate header of the table.

Note
The `.classifier.class` data is the class assigned to the message when pattern database is used. For details, see *Chapter 14, Classifying messages with pattern databases (p. 211)*. The pattern databases provided by BalaBit currently use the following message classes by default: `system`, `security`, `violation`, or `unknown`.

12.3.1. Displaying log statistics

You can choose from **Bar chart** or **Pie chart and List**.

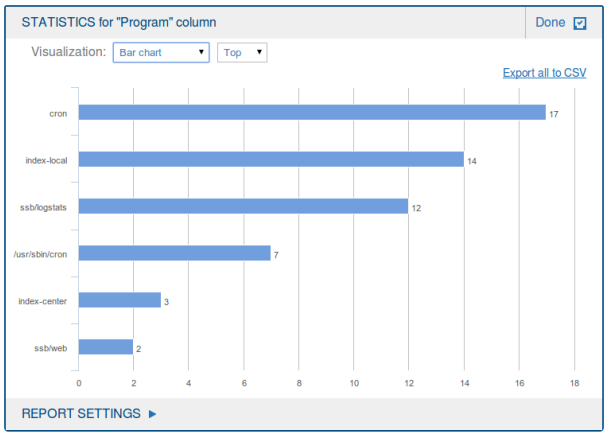

Figure 12.10. Displaying log statistics as Bar chart

In **Pie chart and List** view, percentages add up to 100%. The only exception to this is when statistics are based on **Tags**. Since it provides statistics for tags rather than messages, it is possible that if messages have multiple tags, the percentages will add up to more than 100%.

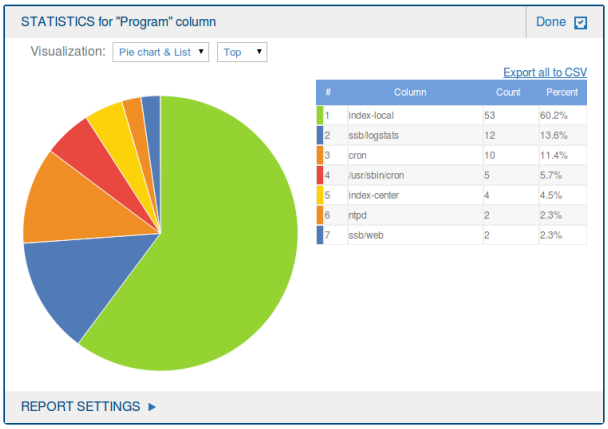

Figure 12.11. Displaying log statistics as Pie chart and List

The statistics start with the largest number of entries. To start statistics with the least number of entries, select **Least**.

You can export these statistics in CSV format using the **Export all to CSV** option, or you can include them in reports as a subchapter.

Warning
Do not use **Export as CSV** to export large amounts of data, as exporting data can be very slow, especially if the system is under heavy load. If you regularly need a large portion of your data in plain text format, consider using the SSB RPC API (for details, see *Chapter 15, The SSB RPC API (p. 221)*), or sharing the log files on the network and processing them with external tools (for details, see *Section 8.6, Accessing log files across the network (p. 150)*).

12.3.2. Procedure – Creating reports from custom statistics

You can save log statistics to include them in reports as a subchapter.

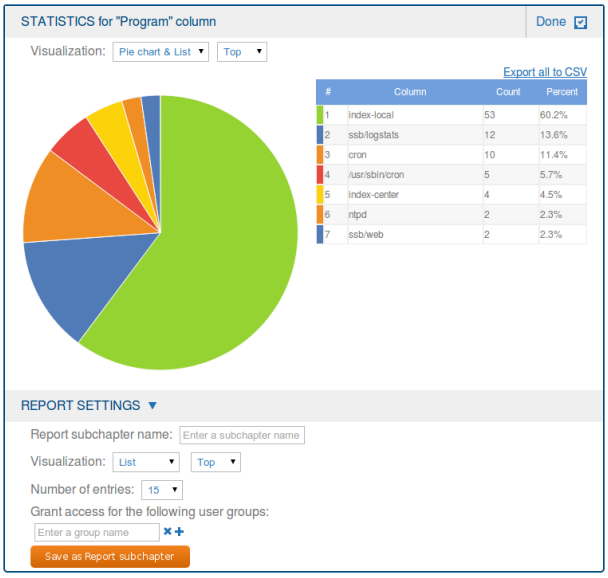
STATISTICS for "Program" column
Done
Visualization: Pie chart & List Top
Export all to CSV
Column Count Percent
1 index-local 53 60.2%
2 ssb/logstats 12 13.6%
3 cron 10 11.4%
4 /usr/sbin/cron 5 5.7%
5 index-center 4 4.5%
6 ntpd 2 2.3%
7 ssb/web 2 2.3%
REPORT SETTINGS
Report subchapter name: Enter a subchapter name
Visualization: List Top
Number of entries: 15
Grant access for the following user groups:
Enter a group name
Save as Report subchapter

Figure 12.12. Creating reports from custom log statistics

Step 1. In the **Statistics** view, click on **Report settings**.

Step 2. Add a name for the statistics in the **Report subchapter name** field.

Step 3. Select the **Visualization** for the report: List, Pie chart, or Bar chart.

Step 4. Choose how the entries are sorted: descending (**Top**) or ascending (**Least**).

Step 5. Choose the **Number of entries** to include.

Note
Selecting **All** includes only the first 1000 results. The remaining results are aggregated as 'others'.

Step 6. Select the user group that can access the subchapter in the **Grant access for the following user groups** field.

Step 7. Click **Save As Report subchapter**.

Step 8. To add the saved subchapter to a report, follow the instructions provided at *Procedure 13.7.3, Configuring custom reports (p. 208)*.

Chapter 13. Browsing the internal messages of SSB

SSB has a separate search interface where you can search, filter, and export internal messages. These internal messages contain the logs created by SSB itself (not the messages collected from external sources), including log messages of the SSB appliance, configuration changes, notifications, alerts, and statistics.

For more information on the internal search interface, see *Section 13.1, Using the internal search interface (p. 199)*.

Log messages of the SSB appliance:

- All available log messages are listed in the **local** logspace in **Search > Spaces**.

  For detailed instructions on using the log search interface, see *Section 12.1, Using the search interface (p. 181)*.

- Recent log messages are also available in **Basic settings > Troubleshooting**.

  For detailed instructions on using the troubleshooting tools, see *Chapter 16, Troubleshooting SSB (p. 222)*.

Configuration changes:

- The configuration-related activity of the SSB users and administrators is available at **AAA > Accounting**. The configuration changes performed on the SSB web interface are all listed here.

  For the list of displayed parameters, see *Section 13.2, Changelogs of SSB (p. 201)*.

- Peers (client computers) that use syslog-ng Premium Edition 3.0 or newer send a special log message to SSB when their configuration is modified. These changes are listed at **Search > Peer configuration change**.

  For the list of displayed parameters, see *Section 13.3, Configuration changes of syslog-ng peers (p. 202)*.

Alerts and notifications:

- If you use the pattern database of SSB to alert on certain log messages, then a history of the alerts is available at **Search > Alerts**.

  For the list of displayed parameters, see *Section 13.4, Log message alerts (p. 203)*.

- Backup and archive notifications, including errors encountered during backup or archiving, are stored at **Search > Archive & Cleanup**.

  For the list of displayed parameters, see *Section 13.5, Notifications on archiving and backups (p. 204)*.

Statistics and reports:

- The statistics of SSB are available at **Basic settings > Dashboard**.

  For the list of available options, see *Section 13.6, Statistics collection options (p. 204)*.

- PDF reports about the configuration changes, system health parameters, and other activities of SSB are available at **Reporting > Reports**.

  For the list of displayed parameters, see *Section 13.7, Reports (p. 205)*.

13.1. Using the internal search interface

The internal search interface is for browsing and filtering the configuration changes, alerts, notifications, and reports of SSB.

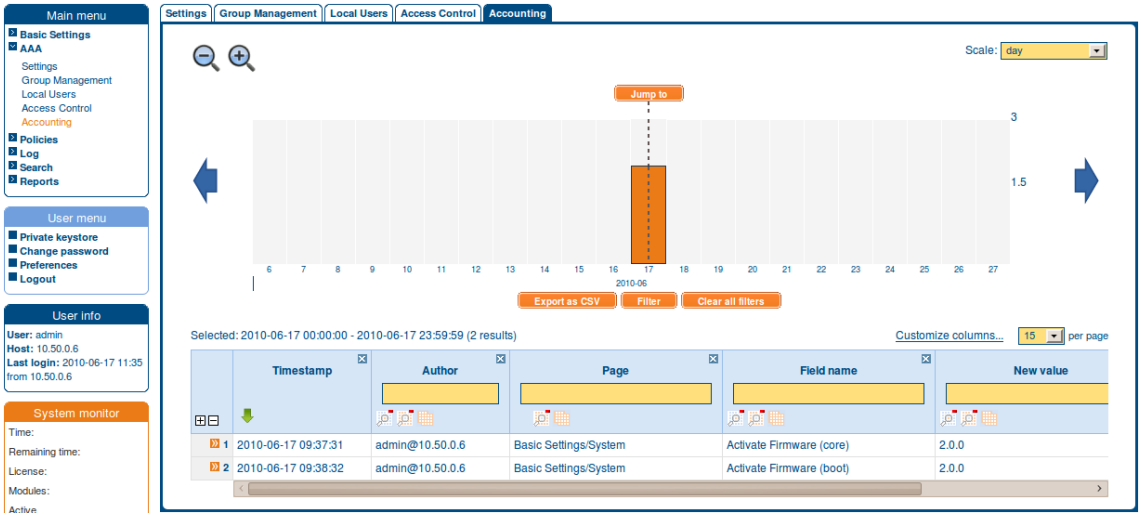

Figure 13.1. The internal search interface

The bars display the number of log messages in the selected interval. Use the icons to zoom, and the arrows to display the previous or the next intervals. To explicitly select a date, select **Jump to** and set the date in the calendar. You can change the length of the displayed interval with the **Scale** option.

Hovering the mouse above a bar displays the number of entries and the start and end date of the period that the bar represents. Click a bar to display the entries of that period in the table. Use Shift+Click to select multiple bars.

If data is too long to fit on one line, it is automatically wrapped and only the first line is displayed. To expand a row, click ⊞. To shrink the row back to its original size, click ⊟. To expand/shrink all rows, click the respective button on the header of the table. The rows can also be expanded/shrunk by double clicking on the respective row.

13.1.1. Filtering

The tables can be filtered for any parameter, or a combination of parameters. To filter the list, enter the filter expression in the input field of the appropriate column, and press **Enter**, or click on an entry in the table.

Note
When you use filters, the bars display the statistics of the filtered results.

Filtering displays also partial matches. For example, filtering the **Author** column on the **AAA > Accounting** screen for `adm` displays all changes performed by users whose username contains the `adm` string.

You can use the icon to perform an exact search, and the icon for inverse filtering ("does not include"). To clear filters from a column, click .

To restore the original table, click **Clear all filters**.

13.1.2. Exporting the results

To save the table of search results as a file, click **Export as CSV**. This saves the table as a text file containing comma-separated values. Note that if an error occurs when exporting the data, the exported CSV file will include a line (usually as the last line of the file) starting with a zero and the details of the problem, for example `0;description_of_the_error`.

Warning
Do not use **Export as CSV** to export large amounts of data, as exporting data can be very slow, especially if the system is under heavy load. If you regularly need a large portion of your data in plain text format, consider using the SSB RPC API (for details, see *Chapter 15, The SSB RPC API (p. 221)*), or sharing the log files on the network and processing them with external tools (for details, see *Section 8.6, Accessing log files across the network (p. 150)*).

13.1.3. Procedure – Customizing columns of the internal search interface

Purpose:

To customize the data displayed on the interface.

Steps:

Step 1. Navigate to the database you want to browse, for example **AAA > Accounting**.

Step 2. Click **Customize Columns**. A pop-up window containing the list of visible and available columns is displayed.

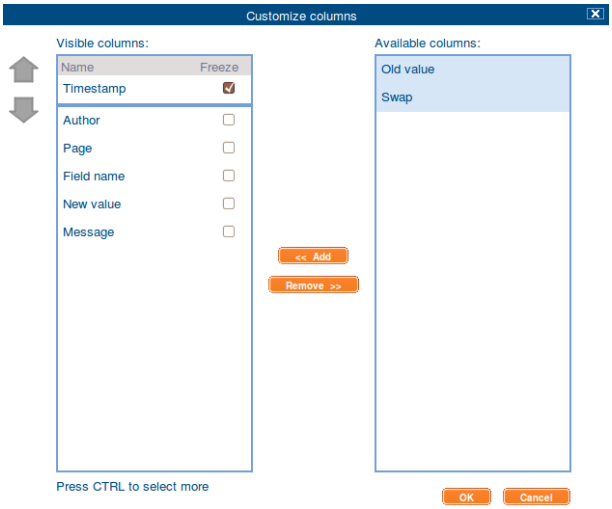

Figure 13.2. Customizing columns of the general search interfaces

Step 3. The displayed parameters are enlisted in the **Visible columns** field. All other available parameters are enlisted in the **Available columns** field.

- To add parameters to the **Visible columns** field, select the desired parameter(s) and click **Add**.
- To remove parameters from the **Visible columns** field, select the desired parameter(s) and click **Remove**.
- To freeze columns (to make them permanently visible, even when scrolling horizontally), enable the **Freeze** option next to the desired parameter.

Note
To select multiple parameters, press **Ctrl** while clicking the items.

Step 4. Click **OK**. The selected information is displayed.

13.2. Changelogs of SSB

SSB automatically records the activity of its users and administrators. These activities are displayed at **AAA > Accounting**. The following information is available:

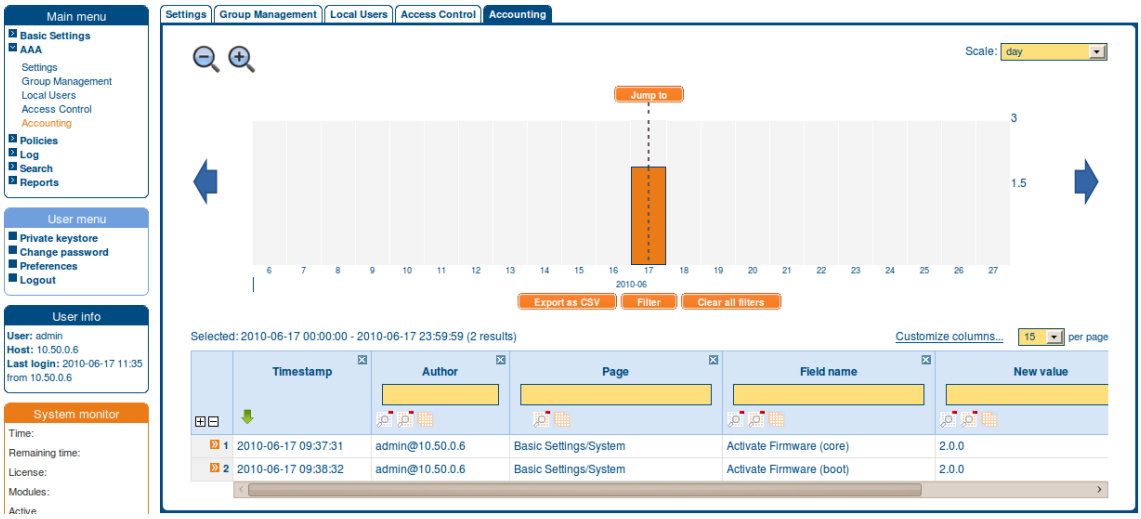

Figure 13.3. Displaying configuration changes

- **Timestamp**: The date when the modification was committed in *YEAR-MONTH-DAY HOUR:MINUTE:SECOND* format.
- **Author**: The SSB user who performed the modification.
- **Page**: The main menu item that was modified (for example *Basic Settings > Management*).
- **Field name**: The name of the field on the page that was modified.
- **New value**: The new value of the field after the modification.
- **Description**: The changelog entered by the SSB administrator. Changelogs are available only if the **AAA > Settings > Require commit log** option was enabled at the time of the change.
- **Old value**: The original value of the field.
- **Swap**: Signs if the order of objects was modified on the page (for example the order of two policies in the list).

13.3. Configuration changes of syslog-ng peers

Peers running syslog-ng Premium Edition 3.0 or later automatically send a notification to SSB when their configuration has changed since the last configuration reload or restart. These log messages are available at **Search > Peer Configuration Change**. Note that the log messages do not contain the actual modification; only indicate that the configuration was modified. The following information is available:

- **Timestamp**: The timestamp received in the message — the time when the log message was created in *YEAR-MONTH-DAY HOUR:MINUTE:SECOND* format.
- **Hostname**: The hostname or IP address of the client whose configuration has been changed.

- **Validity**: The validation of the checksum signature.
- **Version**: Version number of the syslog-ng application that sent the message.
- **Sender address**: The IP address of the client or relay that sent the message directly to SSB.
- **Signature**: The signature of the syslog-ng client.
- **Fingerprint**: The SHA-1 hash of the new configuration file.

13.4. Log message alerts

When using the pattern database, SSB raises alerts for messages that are classified as *`Violation`*. The history of these alerts is available at **Search > Alerts**. The following information is available about the alerts:

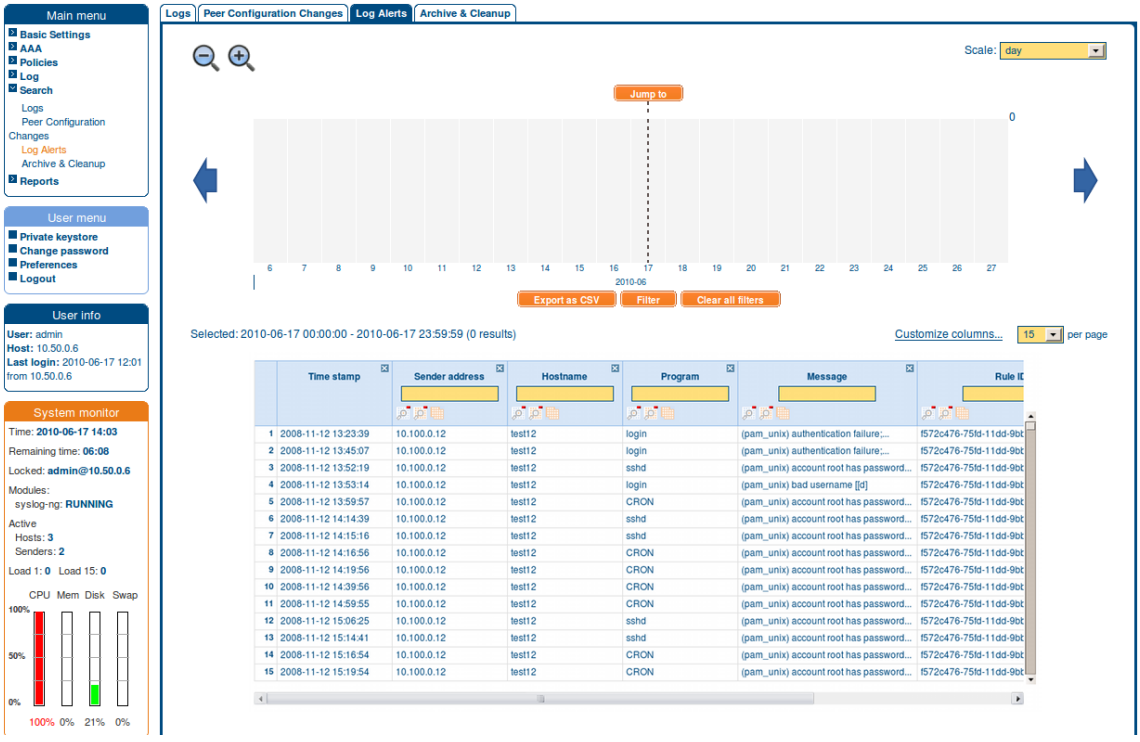

Figure 13.4. Displaying alert messages

- **Timestamp**: The date of the alert in *`YEAR-MONTH-DAY HOUR:MINUTE:SECOND`* format.
- **Sender address**: The IP address of the client or relay that sent the message directly to SSB.
- **Hostname**: The hostname or IP address of the client that sent the message.
- **Program**: The application that generated the message.

- **Message**: The content of the message.
- **Rule ID**: The ID of the classification rule in the pattern database that matched the message. For details, see *Chapter 14, Classifying messages with pattern databases (p. 211)*.
- **Rule description**: The description of the classification rule that matched the message. For details, see *Chapter 14, Classifying messages with pattern databases (p. 211)*.

13.5. Notifications on archiving and backups

Notifications and error messages of the archiving, cleanup and backup procedures are available at **Search** > **Archive & Cleanup**. The following information is available:

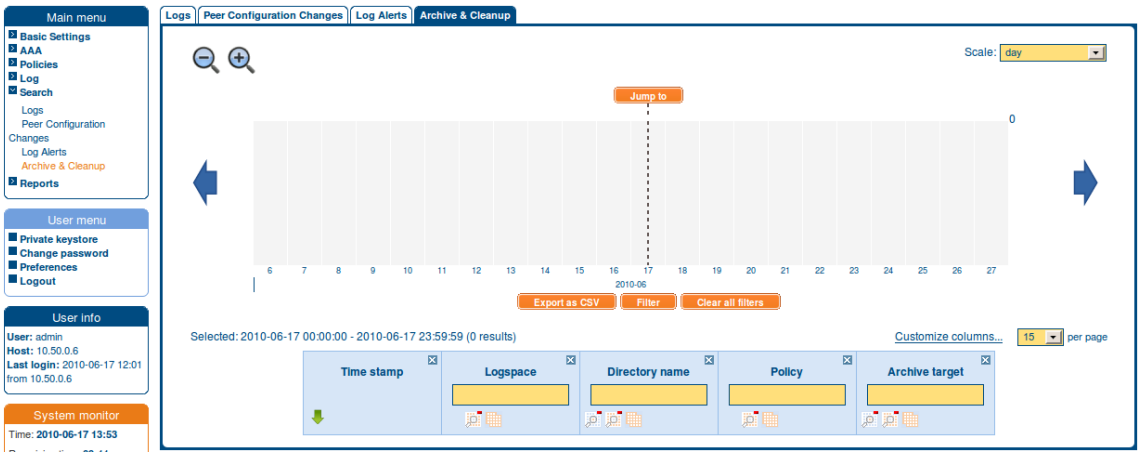

Figure 13.5. Displaying archiving and backup notifications

- **Timestamp**: The date of the message in in `YEAR-MONTH-DAY HOUR:MINUTE:SECOND` format.
- **Logspace**: Name of the archived or backed up logspace.
- **Directory name**: Name of the folder where the archives and backups are located. A new folder is created each day, using the current date as the folder name.
- **Policy**: Name of the archive or backup policy used.
- **Archive target**: Address of the remote server used in the policy.
- **Manual archiving**: Indicates if the archiving or backup process was started manually.

13.6. Statistics collection options

To control the quantity and quality of the statistics collected to the **Dashboard** (for details, see *Section 16.5, Status history and statistics (p. 226)*, set the statistics collection options.

Navigate to **Log > Options > Dashboard Statistics**.

Time-based statistics: the default setting is **Enabled**.

- **Cleanup if unchanged for**: Statistics unchanged (not present in syslog-ng statistics output any more) for this number of days will be cleaned up from the system. Enter `0` here to keep them forever. To start the cleanup process immediately, click **Cleanup now**.

- **Enable statistics for**: the default setting is that all checkboxes are enabled. This allows you to select which options to collect statistics for. To display the collected statistics for an option, navigate to **Basic Settings > Dashboard > Syslog-ng statistics**, select **Time-based statistics** and select the desired option.

Note
When disabling an option, the data will only be deleted after the first cleanup. Until then, the already collected data is still accessible on the dashboard.

Top/Least statistics: the default setting is **Enabled** and all checkboxes are enabled. This allows you to select which options to collect statistics for. To display the collected statistics for an option, navigate to **Basic Settings > Dashboard > Syslog-ng statistics**, select **Top/Least statistics** and select the desired option.

Maximum number of statistics to process: Enter the number of statistics files to keep on the system. Enter `0` here to store unlimited number of statistics files. Statistics over this limit will be dropped, and SSB sends an error message containing the number of entries dropped and the first dropped entry. This setting needs to be increased only if you have more than 10000 hosts.

Sampling interval: Select the sampling interval for the statistics here. A more frequent sampling interval results in more precise graphs at the cost of heavier system load. The default setting is `5 minutes`. The possible parameters are `5 minutes`, `10 minutes`, `30 minutes`, `60 minutes`, `2 hours`, `4 hours`, `8 hours`, `1 day`.

Warning
Hazard of data loss! When changing the Sampling interval, the already existing statistics are not converted to the new sampling rate, but are deleted.

To clear all statistics, click **Clear all statistics**. It is advised to clear statistics if you have changed the number of the statistics files to keep, or if you have disabled the time-based statistics collection.

13.7. Reports

SSB periodically creates reports on the activity of the administrators, the system-health information of SSB, as well as the processed traffic. These reports are available in Portable Document (PDF) format by selecting **Reports > Generated reports** from the Main Menu. The reports are also sent to the e-mail address set at **Basic**

Settings > Management > Mail settings > Send reports to, unless specified otherwise in the configuration of the report.

To access the reports from the SSB web interface, the user must have the appropriate privileges.

Note
If the **Basic Settings > Management > Mail settings > Send reports to** address is not set, the report is sent to the SSB administrator's e-mail address.

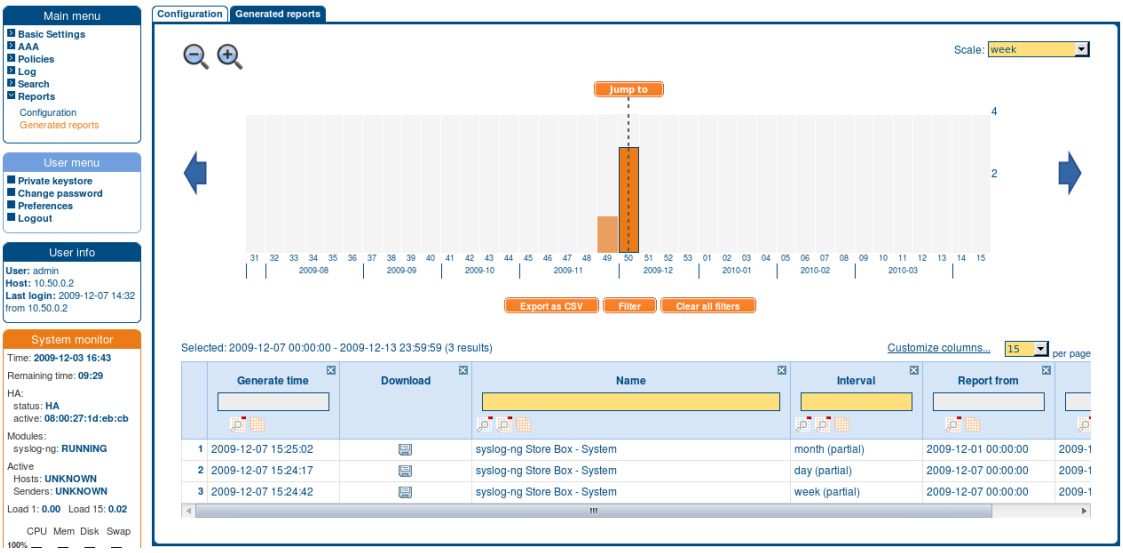

Figure 13.6. Browsing reports

Reports are generated as follows:

- **Daily reports** are generated every day at 00:01.
- **Weekly reports** are generated every week on Monday at 00:01.
- **Monthly reports** are generated on the first day of every month at 00:01.

Tip
Use the time bar to find reports that contain a particular period. If you select a period (for example click on a bar), only those reports will be displayed that contain information about the selected period.

The following information is available about the reports:

- **Download**: A link to download the report.
- **Name**: Name of the report.

- **Interval**: The length of the reported period, for example week, month, and so on.
- **Report from**: The start of the reported interval.
- **Report to**: The end of the reported interval.
- **Generate time**: The date when the report was created.

Tip
To create a report for the current day, select **Generate reports for today**. The report will contain data for the `00:00 - current time` interval. If artificial ignorance (for details, see *Chapter 14, Classifying messages with pattern databases (p. 211)*) is enabled, an artificial ignorance report is created as well.

13.7.1. Contents of the default reports

The default report of SSB (called `System`) is available in Adobe Portable Document Format (PDF), and contains the following information for the given period:

- **Configuration changes**: Lists the number of SSB configuration changes per page and per user. The frequency of the configuration changes is also displayed on a chart.
- **Peer configuration**: Lists the number of times the configuration of a syslog-ng client was changed per client, as well as the version number of the syslog-ng application running on the client (if this information is available).
- **Alerts**: Various statistics about the alerts received from classifying messages using the pattern database (if pattern databases have been uploaded to SSB.
- **syslog-ng traffic statistics**: Displays the rate of incoming, forwarded, stored, and dropped messages in messages/second.
- **System health information**: Displays information about the filesystem and network use of SSB, as well as the average load.

13.7.2. Procedure – Generating partial reports

Purpose:

To generate a report manually for a period that has not been already covered in an automatic report, complete the following steps.

Steps:

Step 1. Login to the SSB web interface, and navigate to **Reports > Configuration**.

Step 2. Select the report you want to generate.

Step 3.

- To create a report from the last daily report till now, click **Generate partial daily report**. For example, if you click this button at 11:30 AM, the report will include the period from 00:01 to 11:30.
- To create a report from the last weekly report till now, click **Generate partial weekly report**. For example, if you click this button on Wednesday at 11:30 AM, the report will include the period from Monday 00:01 to Wednesday 11:30.
- To create a report from the last monthly report till now, click **Generate partial monthly report**. For example, if you click this button at 11:30 AM, December 13, the report will include the period from December 1, 00:01 to December 13, 11:30.

The report will be automatically added in the list of reports (**Reports > Generated reports**), and also sent in an e-mail to the regular recipients of the report.

Step 4. Click

.

13.7.3. Procedure – Configuring custom reports

Purpose:

To configure SSB to create custom reports, complete the following steps with a user that has read & write/perform access to the **use static subchapters** privilege.

Steps:

Step 1. Login to the SSB web interface, and navigate to **Reports > Configuration**.

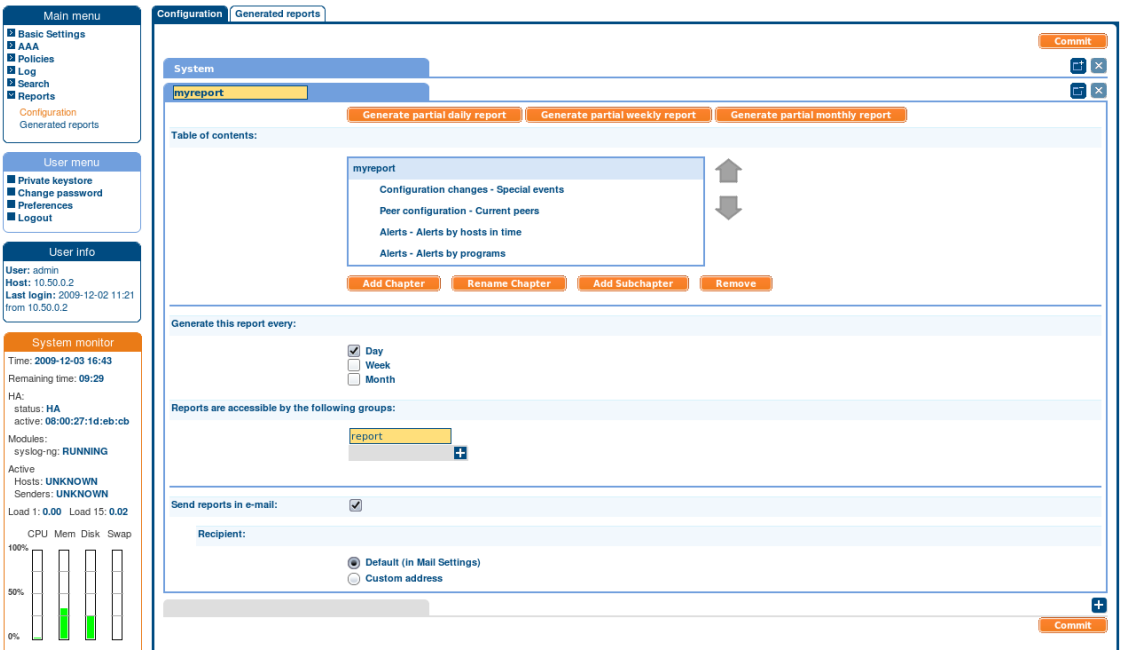

Figure 13.7. Configuring custom reports

Step 2. Click ⊞ and enter a name for the custom report.

Step 3. Reports are organized into chapters and subchapters. Select **Table of contents** > **Add Chapter**, enter a name for the chapter, then click **OK**. Repeat this step to create further chapters if needed.

Step 4. Select **Add Subchapter** to add various reports and statistics to the chapter. The available reports will be displayed in a popup window. The reports created from custom statistics are listed at the end.

Step 5. Use the arrows to change the order of the subchapters if needed.

Step 6. Select how often shall SSB create the report from the **Generate this report every** field. Weekly reports are created on Mondays, while monthly reports on the first day of the month. If you want to generate the report only manually, leave this field empty.

Step 7. By default, members of the *search* group can access the custom reports via the SSB web interface. To change this, enter the name of a different group into the **Reports are accessible by the following groups** field, or click ⊞ to grant access to other groups.

Note
Members of the listed groups will be able to access only these custom reports even if their groups does not have read access to the **Reporting** > **Reports** page. However, only those reports will be listed, to which their group has access to.

Step 8. By default, SSB sends out the reports in e-mail to the address set in the **Basic Settings** > **Management** > **Mail settings** > **Send reports to** field.

Note
If this address is not set, the report is sent to the SSB administrator's e-mail address.

- To disable e-mail sending, unselect the **Send reports in e-mail** option.
- To receive e-mails only when at least one audit trail matching the search criteria was found, unselect the **Send even empty reports** option.
- To e-mail the reports to a different address, select **Recipient > Custom address**, and enter the e-mail address where the reports should be sent. Click ⊕ to list multiple e-mail addresses if needed.

Step 9. Click Commit.

Chapter 14. Classifying messages with pattern databases

Using the pattern database allows you to classify messages into various categories, receive alerts on certain messages, and to collect unknown messages using artificial ignorance.

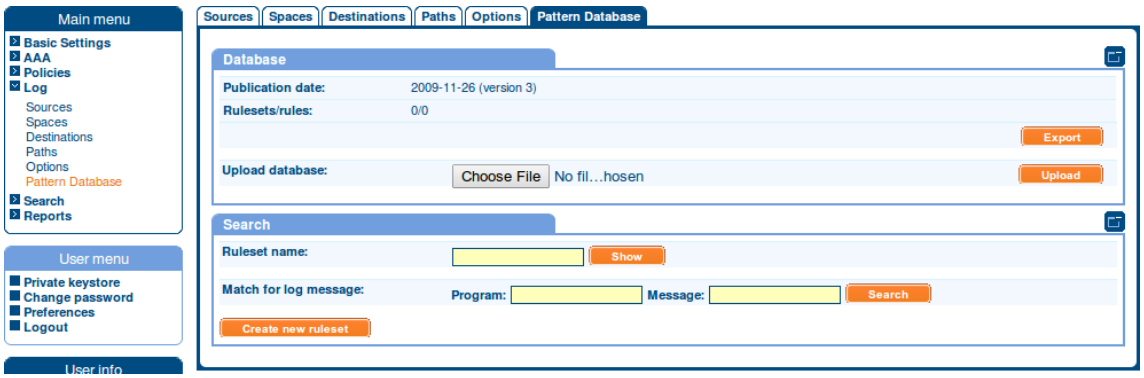

Figure 14.1. Pattern database

Note that the classification of messages is always performed; but its results are used only if you specifically enable the relevant options on the **Log > Options** page.

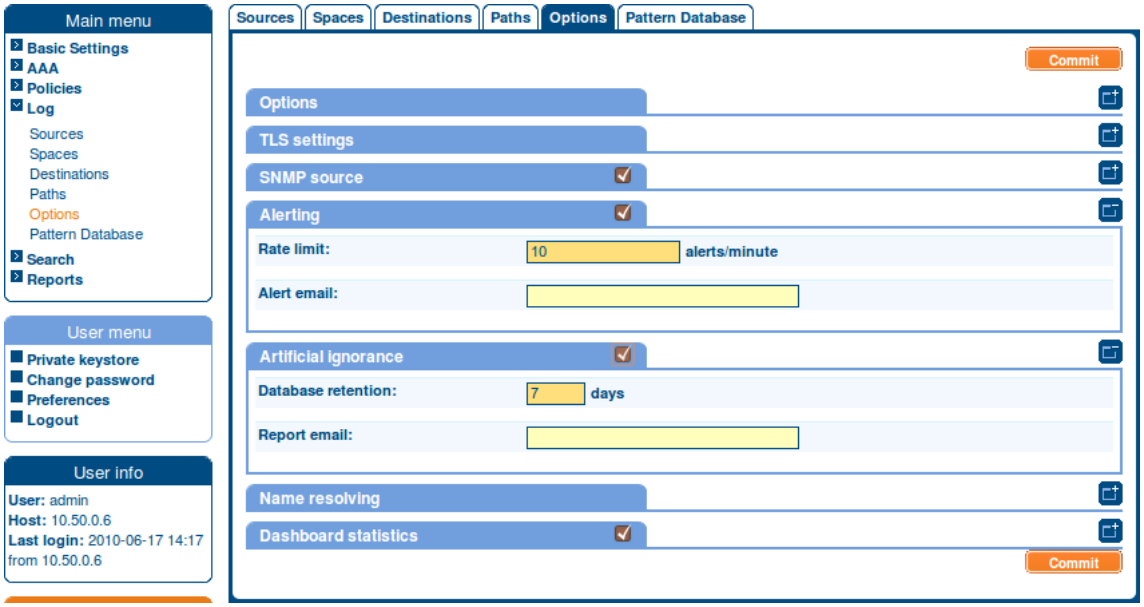

Figure 14.2. Enabling artificial ignorance and pattern-matching alerts

- To receive alerts on messages classified as Violation, navigate to **Log > Options** and enable the **Alerts** option.

- To receive reports on messages not included in the pattern database, navigate to **Log > Options** and enable the **Artificial ignorance** option.

14.1. The structure of the pattern database

The pattern database is organized as follows:

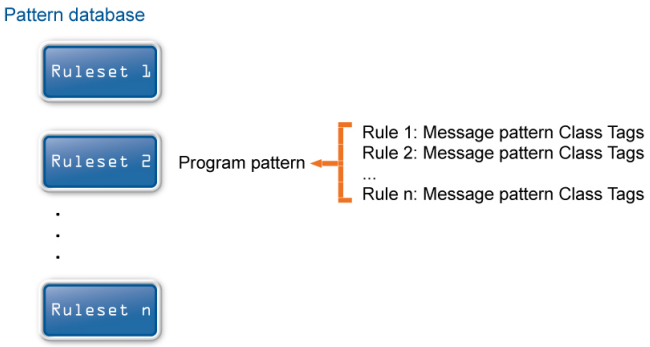

Figure 14.3. The structure of the pattern database

- The pattern database consists of rulesets. A ruleset consists of a Program Pattern and a set of rules: the rules of a ruleset are applied to log messages if the name of the application that sent the message matches the Program Pattern of the ruleset. The name of the application (the content of the ${PROGRAM} macro) is compared to the Program Patterns of the available rulesets, and then the rules of the matching rulesets are applied to the message.
- The Program Pattern can be a string that specifies the name of the appliation or the beginning of its name (for example, to match for sendmail, the program pattern can be sendmail, or just send), and the Program Pattern can contain pattern parsers. Note that pattern parsers are completely independent from the syslog-ng parsers used to segment messages. Additionally, every rule has a unique identifier: if a message matches a rule, the identifier of the rule is stored together with the message.
- Rules consist of a message pattern and a class. The Message Pattern is similar to the Program Pattern, but is applied to the message part of the log message (the content of the ${MESSAGE} macro). If a message pattern matches the message, the class of the rule is assigned to the message (for example, Security, Violation, and so on).
- Rules can also contain additional information about the matching messages, such as the description of the rule, an URL, name-value pairs, or free-form tags. This information is displayed by the syslog-ng Store Box in the e-mail alerts (if alerts are requested for the rule), and are also displayed on the search interface.
- Patterns can consist of literals (keywords, or rather, keycharacters) and pattern parsers.

Note
If the ${PROGRAM} part of a message is empty, rules with an empty Program Pattern are used to classify the message.

If the same Program Pattern is used in multiple rulesets, the rules of these rulesets are merged, and every rule is used to classify the message. Note that message patterns must be unique within the merged rulesets, but the currently only one ruleset is checked for uniqueness.

14.2. How pattern matching works

Figure 14.4. Applying patterns

The followings describe how patterns work. This information applies to program patterns and message patterns alike, even though message patterns are used to illustrate the procedure.

Patterns can consist of literals (keywords, or rather, keycharacters) and pattern parsers. Pattern parsers attempt to parse a sequence of characters according to certain rules.

Note
Wildcards and regular expressions cannot be used in patterns. The `@` character must be escaped, that is, to match for this character, you have to write `@@` in your pattern. This is required because pattern parsers of syslog-ng are enclosed between `@` characters.

When a new message arrives, syslog-ng attempts to classify it using the pattern database. The available patterns are organized alphabetically into a tree, and syslog-ng inspects the message character-by-character, starting from the beginning. This approach ensures that only a small subset of the rules must be evaluated at any given step, resulting in high processing speed. Note that the speed of classifying messages is practically independent from the total number of rules.

For example, if the message begins with the `Apple` string, only patterns beginning with the character `A` are considered. In the next step, syslog-ng selects the patterns that start with `Ap`, and so on, until there is no more specific pattern left.

Note that literal matches take precedence over pattern parser matches: if at a step there is a pattern that matches the next character with a literal, and another pattern that would match it with a parser, the pattern with the literal match is selected. Using the previous example, if at the third step there is the literal pattern `Apport` and a pattern parser `Ap@STRING@`, the `Apport` pattern is matched. If the literal does not match the incoming string (foe example, `Apple`), syslog-ng attempts to match the pattern with the parser. However, if there are two or more parsers on the same level, only the first one will be applied, even if it does not perfectly match the message.

If there are two parsers at the same level (for example, `Ap@STRING@` and `Ap@QSTRING@`), it is random which pattern is applied (technically, the one that is loaded first). However, if the selected parser cannot parse at least one character of the message, the other parser is used. But having two different parsers at the same level is extremely rare, so the impact of this limitation is much less than it appears.

14.3. Searching for rulesets

To display the rules of a ruleset, enter the name of the ruleset into the **Search** > **Ruleset name** field, and click **Show**. If you do not know the name of the ruleset, type the beginning letter(s) of the name, and the names of the matching rulesets will be displayed. If you are looking for a specific rule, enter a search term into the **Program** or **Message** field and select **Search**. The rulesets that contain matching rules will be displayed.

Note
Rulesets containing large number of rules may not display correctly.

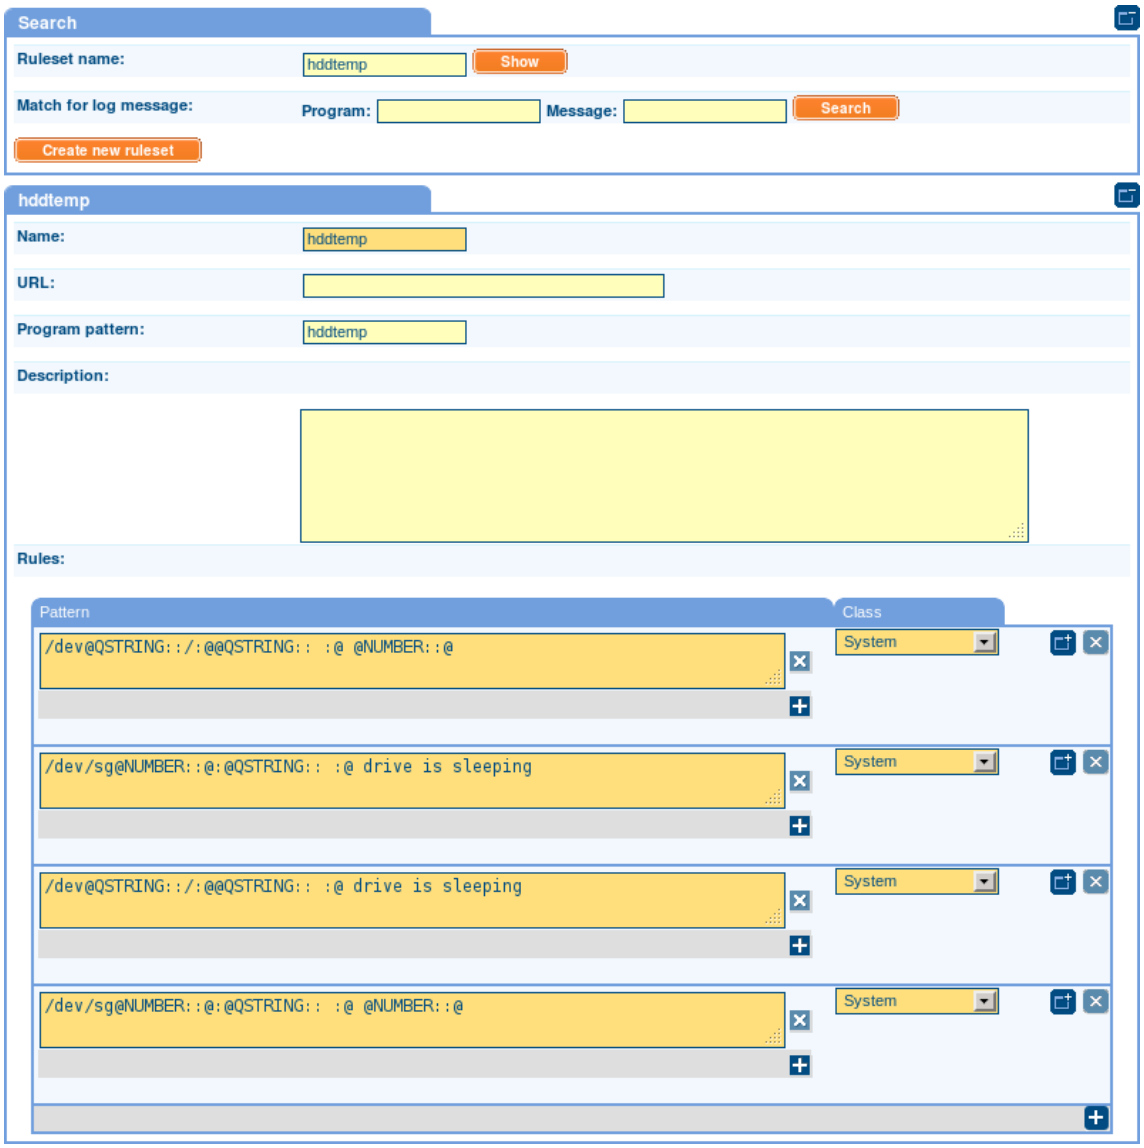

Figure 14.5. Searching rules

14.4. Procedure – Creating new rulesets and rules

Purpose:

To create a new ruleset and new rules, complete the following steps:

Steps:

Step 1. Select **Log > Pattern Database > Create new ruleset**.

Tip
If you search for a ruleset that does not exist, SSB offers you to create a new ruleset with the name you were searching for.

Step 2. Enter a name for the ruleset into the **Name** field.

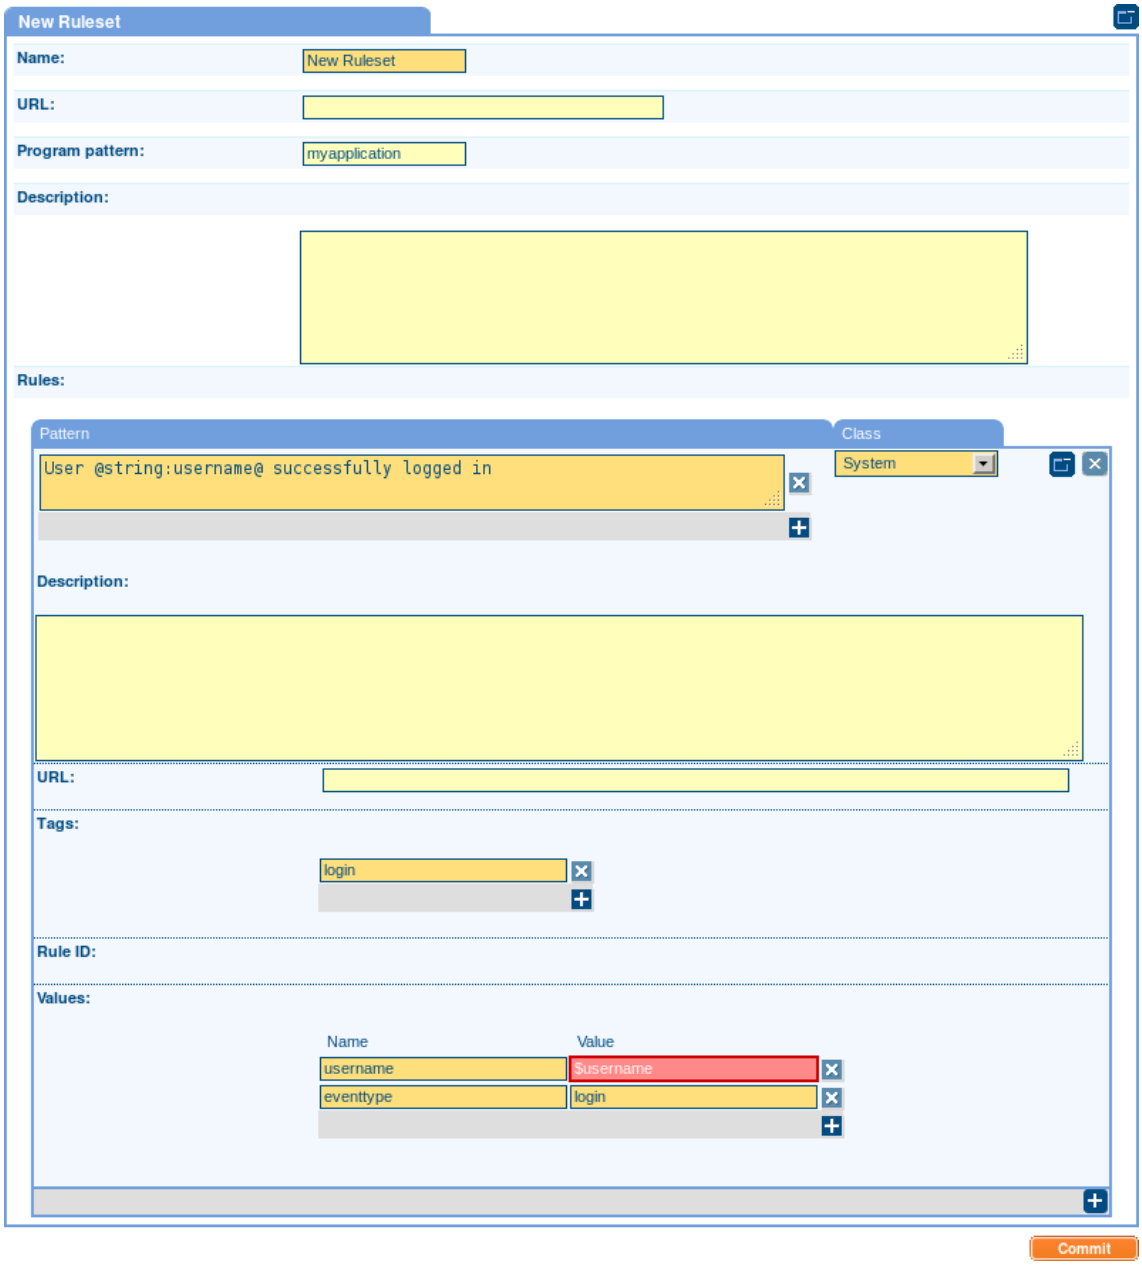

Figure 14.6. Creating pattern database rulesets

Step 3. Enter the name of the application or a pattern that matches the applications into the **Program pattern** field. For details, see *Section 14.7, Using pattern parsers (p. 217)*.

Step 4. Optionally, add a description to the ruleset.

Step 5. Add rules to the class.

Step a. Click ➕ in the **Rules** section.

Step b. Enter the beginning of the log message or a pattern that matches the log message into the **Pattern** field. For details, see *Section 14.7, Using pattern parsers (p. 217)*. Note that only messages sent by applications matching the **Program pattern** will be affected by this pattern.

Step c. Select the type of the message from the **Class** field. This class will be assigned to messages matching the pattern of this rule. The following classes are available: `Violation`, `Security`, and `System`.

If alerting is enabled at **Log > Options > Alerting**, SSB automatically sends an alert if a message is classified as Violation.

Step d. Optionally, you can add a description, custom tags, and name-value pairs to the rule. Note that the values of name-value pairs can contain macros in the `${macroname}` format. For details on pattern databases and macros, see *The syslog-ng Premium Edition Administrator Guide*, available at the *BalaBit Documentation Page*.

Step 6. Repeat the previous step to add more rules.

Step 7. Click Commit.

14.5. Exporting databases and rulesets

To export the entire pattern database, navigate to **Log > Pattern Database** and select **Export**.

To export a ruleset, enter the name of the ruleset into the **Search > Ruleset name** field, click **Show**, and select **Export ruleset**. If you do not know the name of the ruleset, enter a search term into the **Program** or **Message** field and select **Search**. The rulesets that contain matching rules will be displayed.

14.6. Importing pattern databases

You can upload official databases distributed by BalaBit or pattern databases that you have exported from SSB. The following official databases are available at the *BalaBit website*:

- Cisco PIX messages
- Logcheck converted messages
- Zorp messages

To import a ruleset, navigate to **Log > Pattern Database** and select **Browse**. Then locate the database file to upload, and click **Upload**.

Note
Imported rules are effective immediately after the upload is finished.

If you have modified a rule that was originally part of an official database, then the update will not modify this rule.

14.7. Using pattern parsers

Pattern parsers attempt to parse a part of the message using rules specific to the type of the parser. Parsers are enclosed between @ characters. The syntax of parsers is the following:

- a beginning *@* character;
- the type of the parser written in capitals;
- optionally a name;
- parameters of the parser, if any;
- a closing *@* character.

Example 14.1. Pattern parser syntax
A simple parser:

```
@STRING@
```

A named parser:

```
@STRING:myparser_name@
```

A named parser with a parameter:

```
@STRING:myparser_name:*@
```

A parser with a parameter, but without a name:

```
@STRING::*@
```

The following parsers are available:

- *@ANYSTRING@*: Parses everything to the end of the message; you can use it to collect everything that is not parsed specifically to a single macro. In that sense its behavior is similar to the `greedy()` option of the CSV parser.
- *@DOUBLE@*: An obsolete alias of the `@FLOAT@` parser.
- *@ESTRING@*: This parser has a required parameter that acts as the stopcharacter: the parser parses everything until it find the stopcharacter. For example to stop by the next `"` (double quote) character, use `@ESTRING:"@`. As of syslog-ng 3.1, it is possible to specify a stopstring instead of a single character, for example `@ESTRING::stop_here.@`.
- *@FLOAT@*: A floating-point number that may contain a dot (.) character. (Up to syslog-ng 3.1, the name of this parser was `@DOUBLE@`.)
- *@IPv4@*: Parses an IPv4 IP address (numbers separated with a maximum of 3 dots).
- *@IPv6@*: Parses any valid IPv6 IP address.
- *@IPvANY@*: Parses any IP address.

- *@NUMBER@*: A sequence of decimal (0-9) numbers (for example 1, 0687, and so on). Note that if the number starts with the 0x characters, it is parsed as a hexadecimal number, but only if at least one valid character follows 0x.
- *@QSTRING@*: Parse a string between the quote characters specified as parameter. Note that the quote character can be different at the beginning and the end of the quote, for example: `@QSTRING:"@` parses everything between two quotation marks (`"`), while `@QSTRING:<>@` parses from an opening bracket to the closing bracket.
- *@STRING@*: A sequence of alphanumeric characters (0-9, A-z), not including any whitespace. Optionally, other accepted characters can be listed as parameters (for example to parse a complete sentence, add the whitespace as parameter, like: `@STRING: @`). Note that the `@` character cannot be a parameter, nor can line-breaks or tabs.

Patterns and literals can be mixed together. For example, to parse a message that begins with the `Host:` string followed by an IP address (for example `Host: 192.168.1.1`), the following pattern can be used: `Host:@IPv4@`.

Note
Note that using parsers is a CPU-intensive operation. Use the ESTRING and QSTRING parsers whenever possible, as these can be processed much faster than the other parsers.

Example 14.2. Using the STRING and ESTRING parsers
For example, if the message is `user=joe96 group=somegroup`, `@STRING:mytext:@` parses only to the first non-alphanumeric character (=), parsing only `user`. `@STRING:mytext:=@` parses the equation mark as well, and proceeds to the next non-alphanumeric character (the whitespace), resulting in `user=joe96` being parsed. `@STRING:mytext:=@` will parse the whitespace as well, and proceed to the next non-alphanumeric non-equation mark non-whitespace character, resulting in `user=joe96 group=somegroup`.

Of course, usually it is better to parse the different values separately, like this: `"user=@STRING:user@ group=@STRING:group@"`.

If the username or the group may contain non-alphanumeric characters, you can either include these in the second parameter of the parser (as shown at the beginning of this example), or use an ESTRING parser to parse the message till the next whitespace: `"user=@ESTRING:user: @group=@ESTRING:group: @"`.

Example 14.3. Patterns for multiline messages
Patterns can be created for multiline log messages. For example, the following pattern will find the multiline message where a line ends with `first` and the next line starts with `second`:

```
first
second
```

14.8. Procedure – Using parser results in filters and templates

Purpose:

The results of message classification and parsing can be used in custom filters and file and database templates as well. There are two built-in macros in SSB that allow you to use the results of the classification: the `.classifier.class` macro contains the class assigned to the message (for example violation, security, or unknown), while the `.classifier.rule_id` macro contains the identifier of the message pattern that matched the message.

Note
ID of the message pattern is automatically inserted into the template if the messages are forwarded to an SQL database.

To use these macros as filters in a log path, complete the following procedure:

Steps:

Step 1. Navigate to **Log > Paths** and select the log path to use.

Step 2. To filter on a specific message class, select **Add filter > classifier_class**, select ⊞, then select the class to match (for example *`Violation`*) from the **classifier_class** field.

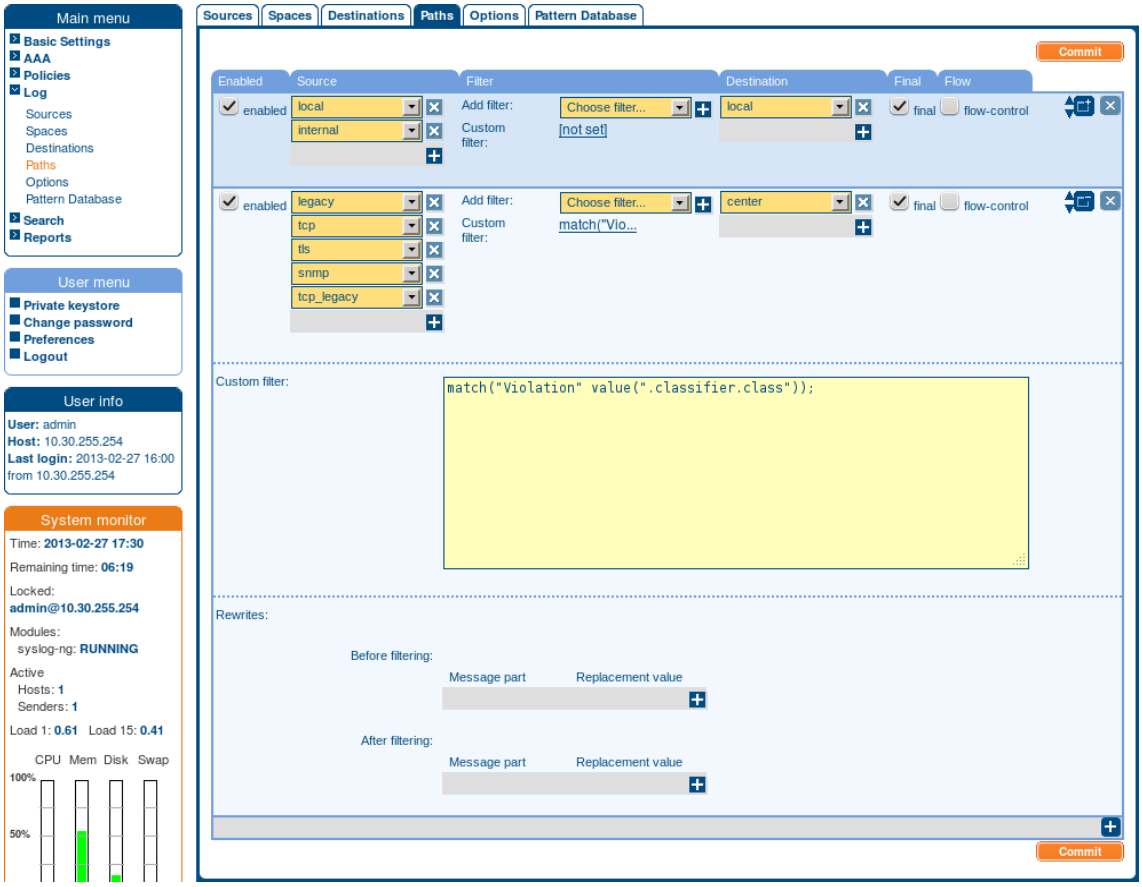

Figure 14.7. Filtering messages based on the classification

Step 3. To filter on messages matching a specific classification rule, **Add filter > classifier_rule_id**, select ⊞, then enter the unique identifier of the rule (for example *`e1e9c0d8-13bb-11de-8293-000c2922ed0a`*) into the **classifier_rule_id** field.

Note
To filter messages based on other classification data like tags, you have to use Custom filters. For details, see *Section 10.3, Filtering messages (p. 169)*.

Step 4. Click Commit.

14.9. Using the values of pattern parsers in filters and templates

Similarly, to *Procedure 14.8, Using parser results in filters and templates (p. 218)*, the results of pattern parsers can be used as well. To accomplish this, you have to add a name to the parser, and then you can use this name as a macro that refers to the parsed value of the message.

For example, you want to parse messages of an application that look like *"Transaction: <type>."*, where *<type>* is a string that has different values (for example refused, accepted, incomplete, and so on). To parse these messages, you can use the following pattern:

```
'Transaction: @ESTRING::.@'
```

Here the @ESTRING@ parser parses the message until the next full stop character. To use the results in a filter or a filename template, include a name in the parser of the pattern, for example:

```
'Transaction:
            @ESTRING:TRANSACTIONTYPE:.@'
```

After that, add a custom template to the logpath that uses this template. For example, to select every *accepted* transaction, use the following custom filter in the log path:

```
match("accepted" value("TRANSACTIONTYPE"));
```

Note
The above macros can be used in database columns and filename templates as well, if you create custom templates for the destination or logspace.

Chapter 15. The SSB RPC API

Version 3.2 and later of syslog-ng Store Box can be accessed using a Remote-Procedure Call Application Programming Interface (RPC API).

The SSB RPC API allows you to access and query SSB logspaces from remote applications. You can access the API using a RESTful protocol over HTTPS, meaning that you can use any programming language that has access to a RESTful HTTPS client to integrate SSB to your environment. Sample shell code snippets are provided in the API documentation.

Accessing SSB with the RPC API offers several advantages:

- Integration into custom applications and environments
- Flexible, dynamic search queries

15.1. Requirements for using the RPC API

To access SSB using the RPC API, the following requirements must be met:

- The appliance can be accessed using a RESTful protocol over authenticated HTTPS connections.
- The user account used to access SSB via RPC must have **Search** privilege (which provides access to all logspaces), or must be a member of the groups listed in the **Access Control** option of the particular logspace. For details on managing user privileges, see *Procedure 5.6.1, Modifying group privileges (p. 78)*.

15.2. RPC client requirements

The client application used to access SSB must meet the following criteria:

- Support RESTful web APIs over HTTPS
- Properly handle complex object types
- Include a JSON decoder for interpreting the results of search operations

15.3. Documentation of the RPC API

The documentation of the SSB RPC API is available online from the following URL: *https://<ip-address-of-SSB>/api/1/documentation*. This documentation contains the detailed description of public calls, with examples.

Chapter 16. Troubleshooting SSB

This section describes the tools to detect networking problems, and also how to collect core files and view the system logs of SSB.

If you need to find the SSB appliance in the server room, you can use IPMI to control the front panel identify light. On syslog-ng Store Box SSB5000, and SSB10000, navigate to **Basic Settings** > **System** > **Hardware informantion** > **Blink system identification lights** and click **On** to blink the LEDs of hard disk trays on the front of the SSB appliance in red.

16.1. Procedure – Network troubleshooting

Purpose:

The **Troubleshooting** menu provides a number of diagnostic commands to resolve networking issues. Logfiles of SSB can also be displayed here — for details, see *Procedure 16.3, Viewing logs on SSB (p. 223)*.

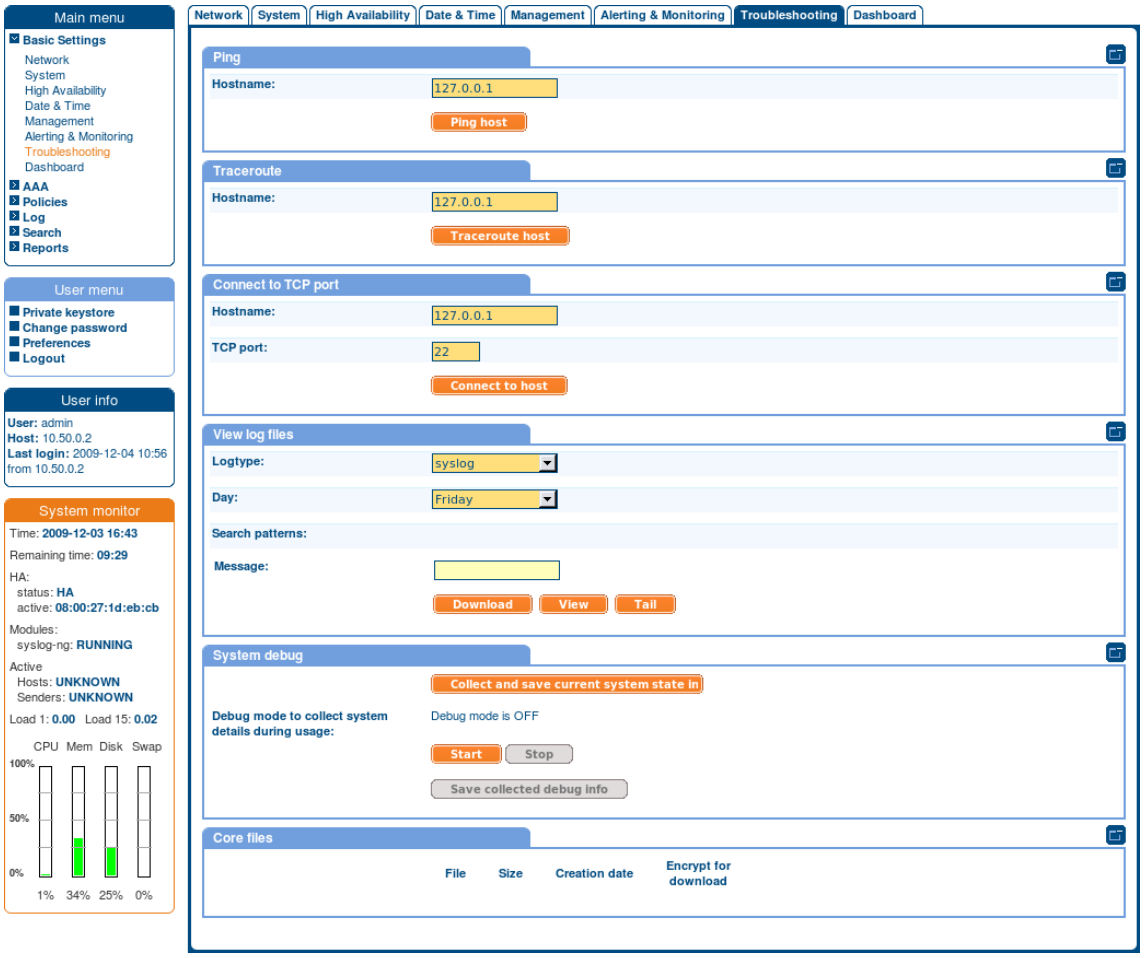

Figure 16.1. Network troubleshooting with SSB

The following commands are available:

- `ping`: Sends a simple message to the specified host to test network connectivity.
- `traceroute`: Sends a simple message from SSB to the specified host and displays all hosts on the path of the message. It is used to trace the path the message travels between the hosts.
- `connect`: Attempts to connect the specified host using the specified port. It is used to test the availability or status of an application on the target host.

To execute one of the above commands, complete the following steps:

Steps:

Step 1. Navigate to **Basic Settings** > **Troubleshooting**.

Step 2. Enter the IP address or the hostname of the target host into the **Hostname** field of the respective command. For the `Connect` command, enter the target port into the **Port** field.

Step 3. Click the respective action button to execute the command.

Step 4. Check the results in the popup window. Log files are displayed in a separate browser window.

16.2. Gathering data about system problems

SSB automatically generates core files if an important software component (for example syslog-ng, or the indexer) of the system crashes for some reason. These core files can be of great help to the BalaBit Support Team to identify problems. When a core file is generated, the SSB administrator receives an alerting e-mail, and an SNMP trap is generated if alerting is properly configured (for details, see *Section 4.6, Configuring system monitoring on SSB (p. 47)* and *Section 4.5, SNMP and e-mail alerts (p. 41)*). To display a list of alerts if monitoring is not configured, navigate to **Search** > **Log Alerts**.

To list and download the generated core files, navigate to **Basic Settings** > **Troubleshooting** > **Core files**.

By default, core files are deleted after 14 days. To change the deletion timeframe, navigate to **Basic Settings** > **Management** > **Core files**.

16.3. Procedure – Viewing logs on SSB

Purpose:

The **Troubleshooting** menu provides an interface to view the logs generated by the various components of SSB. For details on how to browse the log messages received by SSB from its peers, see *Chapter 12, Browsing log messages (p. 181)*.

Note
Because of performance reasons, log files larger than 2 Megabytes are not displayed in the web interface. To access these logs, download the file instead.

Steps:

Step 1. Navigate to **Basic Settings** > **Troubleshooting** > **View log files**.

Step 2. Use the **Logtype** roll-down menu to select the message type.

- *SSB*: Logs of the SSB web interface.
- *syslog*: All system logs of the SSB host.
- *syslog-ng*: Internal log messages of the built-in syslog-ng server. These logs do not contain messages received from the peers.

Step 3.

- To download the log file, click **Download**.
- To follow the current log messages real-time, click **Tail**.
- To display the log messages, click **View**.

Step 4. To display log messages of the last seven days, select the desired day from the **Day:** field and click **View**.

Tip
To display only the messages of a selected host or process, enter the name of the host or process into the **Message:** field.

The **Message:** field acts as a generic filter: enter a keyword or a POSIX (basic) regular expression to display only messages that contain the keyword or match the expression.

16.4. Procedure – Collecting logs and system information for error reporting

Purpose:

To track down support requests, the BalaBit Support Team might request you to collect system-state and debugging information. This information is collected automatically, and contains log files, the configuration file of SSB, and various system-statistics.

Note
Sensitive data like key files and passwords are automatically removed from the files.

The **Basic Settings > Management > Debug logging > Enable debug logs** option is not related to the verbosity of log messages: it adds the commands executed by the SSB web interface to the log.

To collect system-state information, navigate to **Basic Settings > Troubleshooting > System debug** and click **Collect and save current system state info**, then save the created zip file. The name of the file uses the *debug_info-<hostname>YYYYMMDDHHMM* format.

To collect information for a specific error, complete the following steps:

Steps:

Step 1. Navigate to **Basic Settings > Troubleshooting > System debug**.

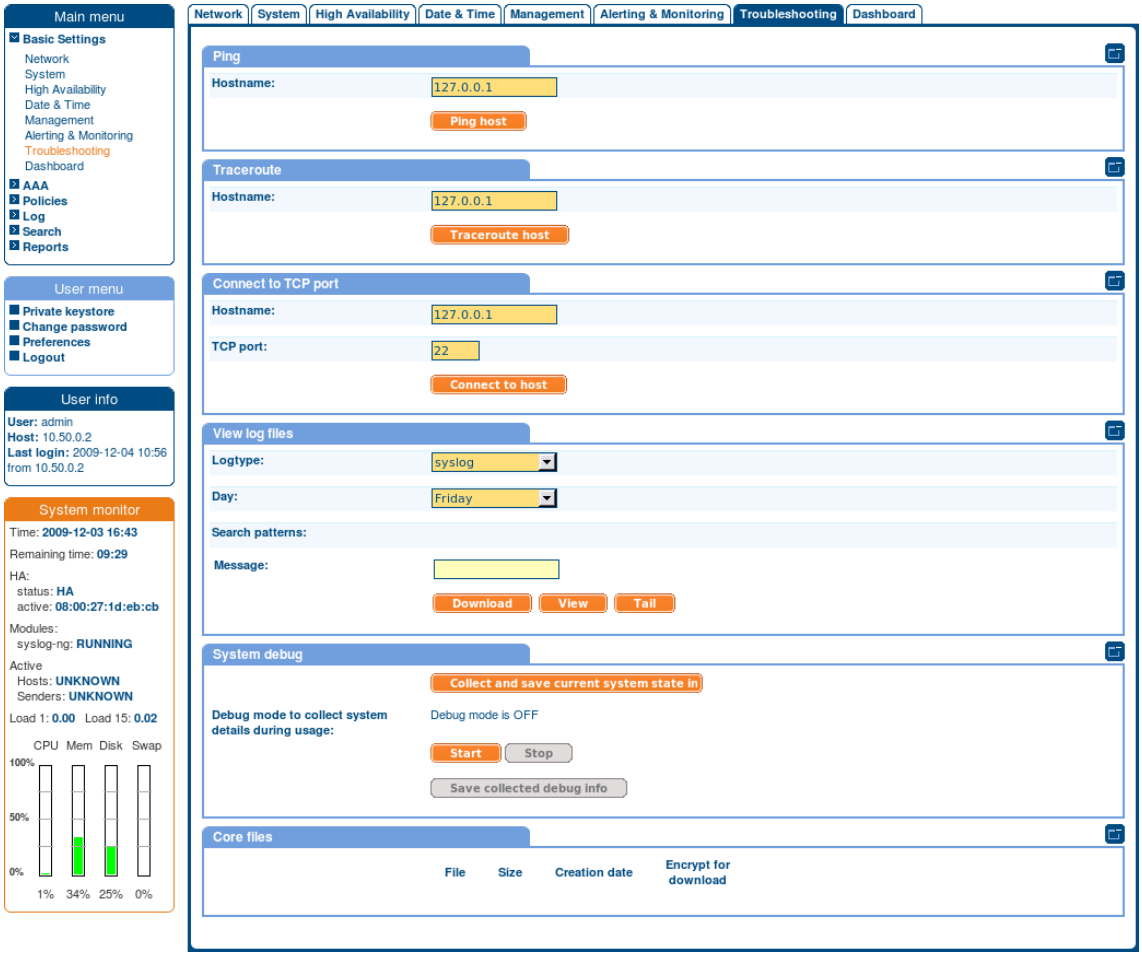

Figure 16.2. Collecting debug information

Step 2. Click **Start**.

Note
Starting debug mode increases the log level of SSB, and might cause performance problems if the system is under a high load.

Step 3. Reproduce the event that causes the error, for example send a log message from a client.

Step 4. Click **Stop**.

Step 5. Click **Save the collected debug info** and save the created zip file. The name of the file uses the *debug_info-<hostname>YYYYMMDDHHMM* format.

Step 6. Attach the file to your support ticket.

16.5. Status history and statistics

SSB displays various statistics and status history of system data and performance on the dashboard at **Basic Settings > Dashboard**. The dashboard is essentially an extension of the system monitor: the system monitor displays only the current values, while the dashboard creates graphs and statistics of the system parameters.

The dashboard consists of different modules. Every module displays the history of a system parameter for the current day. To display the graph for a longer period (last week, last month, or last year), select the **Week**, **Month**, or **Year** options, respectively. Hovering the mouse over a module enlarges the graph and displays the color code used on the graph.

To display statistics of a module as a table for the selected period, click on the graph.

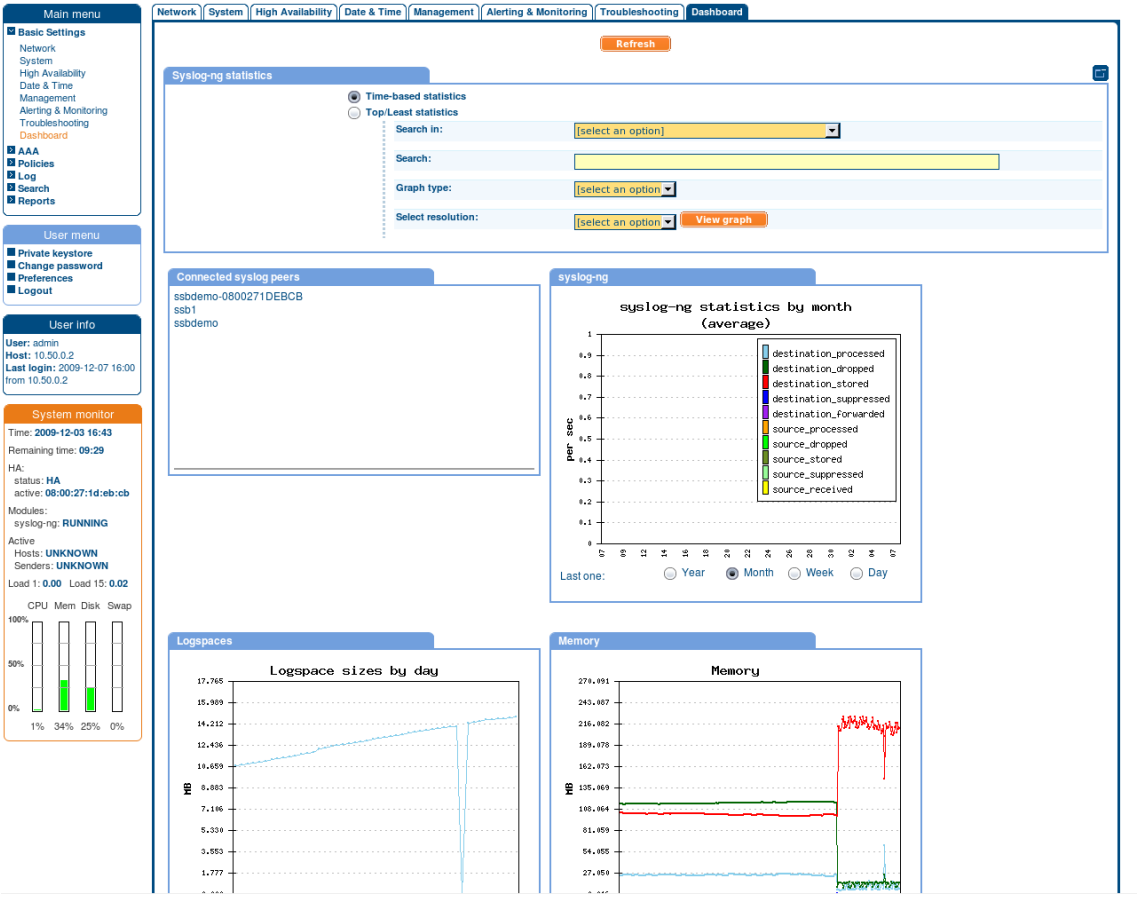

Figure 16.3. The dashboard

The following modules are displayed on the dashboard of SSB:

- **syslog-ng**: syslog-ng statistics about the received, processed, and dropped messages. See also *Procedure 16.5.1, Displaying custom syslog-ng statistics (p. 227)*.

- **Connected syslog peers**: A list of hosts that actively send messages to SSB. Note that these values are updated periodically based on the **Sampling interval** set on page **Log > Options > Dashboard Statistics**. For details, see *Procedure 16.5.1, Displaying custom syslog-ng statistics (p. 227)*.

- **syslog-ng statistics**: The rate of incoming messages in messages/second. Note that the values displayed are an average values calculated for the last fifteen minutes.
- **Logspaces**: The size of the logspaces. Note that these values are updated only in every ten minutes.
- **Memory**: The memory used by the system.
- **Disk**: Filesystem usage for the different partitions.
- **CPU**: CPU usage.
- **Network connections**: Number of network connections.
- **External interface**: Traffic on the external interface.
- **Management interface**: Traffic on the management interface.
- **Load average**: Average load of the system.
- **Processes**: The number of running processes.

For details about setting the statistics collection options, see *Section 13.6, Statistics collection options (p. 204)*

16.5.1. Procedure – Displaying custom syslog-ng statistics

Purpose:

To display statistics of a specific source, destination, or host, complete the following procedure:

Steps:

Step 1. Navigate to **Basic Settings > Dashboard > syslog-ng statistics**.

Step 2.
- To display the statistics of a destination file, select *`destination`* from the **Search in** field, and enter the name of the destination into the **Search** field. Destinations name all start with the *`ds`* characters.
- To display the statistics of a particular host, select *`source`* from the **Search in** field, and enter the hostname or IP address of the host into the **Search** field.

Step 3. Select the time period to display from the **Select resolution** field.

Step 4. Click **View graph**.

16.6. Troubleshooting an SSB cluster

The following sections help you to solve problems related to high availability clusters.

- For a description of the possible statuses of the SSB cluster and its nodes, the DRBD data storage system, and the heartbeat interfaces (if configured), see *Section 16.6.1, Understanding SSB cluster statuses (p. 228)*.
- To recover a cluster that has broken down, see *Procedure 16.6.2, Recovering SSB if both nodes broke down (p. 230)*.
- To resolve a split-bran situation when the nodes of the cluster were simultaneously active for a time, see *Procedure 16.6.3, Recovering from a split brain situation (p. 231)*.
- To replace a broken node with a new appliance, see *Procedure 16.6.4, Replacing a node in an SSB HA cluster (p. 233)*.

16.6.1. Understanding SSB cluster statuses

This section explains the possible statuses of the SSB cluster and its nodes, the DRBD data storage system, and the heartbeat interfaces (if configured). SSB displays this information on the **Basic Settings > High Availability** page.

The **Status** field indicates whether the SSB nodes recognize each other properly and whether those are configured to operate in high availability mode. The status of the individual SSB nodes is indicated in the **Node HA status** field of the each node. The following statuses can occur:

- **Standalone**: There is only one SSB unit running in `standalone` mode, or the units have not been converted to a cluster (the **Node HA status** of both nodes is `standalone`). Click **Convert to Cluster** to enable High Availability mode.
- **HA**: The two SSB nodes are running in High Availability mode. **Node HA status** is *HA* on both nodes, and the **Node HA UUID** is the same on both nodes.
- **Half**: High Availability mode is not configured properly, one node is in `standalone`, the other one in *HA* mode. Connect to the node in *HA* mode, and click **Join HA** to enable High Availability mode.
- **Broken**: The two SSB nodes are running in High Availability mode. **Node HA status** is *HA* on both nodes, but the **Node HA UUID** is different. Contact the BalaBit Support Team for help. For contact details, see *Section 5, Contact and support information (p. xiii)*.
- **Degraded**: SSB was running in high availability mode, but one of the nodes has disappeared (for example broken down, or removed from the network). Power on, reconnect, or repair the missing node.
- **Degraded (Disk Failure)**: A hard disk of the slave node is not functioning properly and must be replaced. To request a replacement hard disk and for details on replacing the hard disk, contact the *BalaBit Support Team*.

- **Degraded Sync**: Two SSB units were joined to High Availability mode, and the first-time synchronization of the disks is currently in progress. Wait for the synchronization to complete. Note that in case of large disks with lots of stored data, synchronizing the disks can take several hours.

- **Split brain**: The two nodes lost the connection to each other, with the possibility of both nodes being active (master) for a time.

Warning
Hazard of data loss! In this case, valuable log messages might be available on both SSB nodes, so special care must be taken to avoid data loss. For details on solving this problem, see *Procedure 16.6.3, Recovering from a split brain situation (p. 231)*.

Do NOT reboot or shut down the nodes.

- **Invalidated**: The data on one of the nodes is considered out-of-sync and should be updated with data from the other node. This state usually occurs during the recovery of a split-brain situation when the DRBD is manually invalidated.

- **Converted**: After converting nodes to a cluster (clicking **Convert to Cluster**) or enabling High Availability mode (clicking **Join HA**) and before rebooting the node(s).

Note
If you experience problems because the nodes of the HA cluster do not find each other during system startup, navigate to **Basic Settings > High Availability** and select **Make HA IP permanent**. That way the IP address of the HA interfaces of the nodes will be fix, which helps if the HA connection between the nodes is slow.

The **DRBD status** field indicates whether the latest data (including SSB configuration, log files, and so on) is available on both SSB nodes. The master node (this node) must always be in **consistent** status to prevent data loss. Inconsistent status means that the data on the node is not up-to-date, and should be synchronized from the node having the latest data.

The **DRBD status** field also indicates the connection between the disk system of the SSB nodes. The following statuses are possible:

- **Connected**: Both nodes are functioning properly.

- **Connected (Disk Failure)**: A hard disk of the slave node is not functioning properly and must be replaced. To request a replacement hard disk and for details on replacing the hard disk, contact the *BalaBit Support Team*.

- **Invalidated**: The data on one of the nodes is considered out-of-sync and should be updated with data from the other node. This state usually occurs during the recovery of a split-brain situation when the DRBD is manually invalidated.

- **Sync source** or **Sync target**: One node (**Sync target**) is downloading data from the other node (**Sync source**).

When synchronizing data, the progress and the remaining time is displayed in the **System monitor**.

Warning
When the two nodes are synchronizing data, do not reboot or shutdown the master node. If you absolutely must shutdown the master node during synchronization, shutdown the slave node first, and then the master node.

- **Split brain**: The two nodes lost the connection to each other, with the possibility of both nodes being active (master) for a time.

Warning
Hazard of data loss! In this case, valuable log messages might be available on both SSB nodes, so special care must be taken to avoid data loss. For details on solving this problem, see *Procedure 16.6.3, Recovering from a split brain situation (p. 231)*.

- **WFConnection**: One node is waiting for the other node; the connection between the nodes has not been established yet.

If a redundant heartbeat interface is configured, its status is also displayed in the **Redundant Heartbeat status** field, and also in the **HA > Redundant** field of the System monitor. For a description of redundant heartbeat interfaces, see *Procedure 6.2.3, Redundant heartbeat interfaces (p. 90)*.

The possible status messages are explained below.

- **NOT USED**: There are no redundant heartbeat interfaces configured.
- **OK**: Normal operation, every redundant heartbeat interface is working properly.
- **DEGRADED-WORKING**: Two or more redundant heartbeat interfaces are configured, and at least one of them is functioning properly. This status is displayed also when a new redundant heartbeat interface has been configured, but the nodes of the SSB cluster has not been restarted yet.
- **DEGRADED**: The connection between the redundant heartbeat interfaces has been lost. Investigate the problem to restore the connection.
- **INVALID**: An error occurred with the redundant heartbeat interfaces. Contact the BalaBit Support Team for help. For contact details, see *Section 5, Contact and support information (p. xiii)*.

16.6.2. Procedure – Recovering SSB if both nodes broke down

Purpose:

It can happen that both nodes break down simultaneously (for example because of a power failure), or the slave node breaks down before the original master node recovers. To properly recover SSB, complete the following steps:

Note
As of SSB version 1.1.1, when both nodes of a cluster boot up in parallel, the node with the `1.2.4.1` HA IP address will become the master node.

Steps:

Step 1. Power off both nodes by pressing and releasing the power button.

Warning
Hazard of data loss! If SSB does not shut off, press and hold the power button for approximately 4 seconds. This method terminates connections passing SSB and might result in data loss.

Step 2. Power on the node that was the master before SSB broke down. Consult the system logs to find out which node was the master before the incident: when a node boots as master, or when a takeover occurs, SSB sends a log message identifying the master node.

Tip
Configure remote logging to send the log messages of SSB to a remote server where the messages are available even if the logs stored on SSB become unaccessible. For details on configuring remote logging, see *Section 4.5, SNMP and e-mail alerts (p. 41)*.

Step 3. Wait until this node finishes the boot process.

Step 4. Power on the other node.

16.6.3. Procedure – Recovering from a split brain situation

Purpose:

A split brain situation is caused by a temporary failure of the network link between the cluster nodes, resulting in both nodes switching to the active (master) role while disconnected. This might cause that new data (for example log messages) is created on both nodes without being replicated to the other node. Thus, it is likely in this situation that two diverging sets of data have been created, which cannot be trivially merged.

Warning
Hazard of data loss! In a split brain situation, valuable log messages might be available on both SSB nodes, so special care must be taken to avoid data loss.

The nodes of the SSB cluster automatically recognize the split brain situation once the connection between the nodes is reestablished, and do not perform any data synchronization to prevent data loss. When a split brain situation is detected, it is visible on the SSB system monitor, in the system logs (`Split-Brain detected, dropping connection!`), and SSB sends an alert as well.

To recover an SSB cluster from a split brain situation, complete the following steps.

Warning
Do NOT shut down the nodes.

Steps:

Step 1. Temporarily disable all incoming traffic. Navigate to **Basic Settings** > **System** > **Service control** > **Syslog traffic, indexing & search:** and click **Disable**.

If the web interface is not accessible or unstable, complete the following steps:

Step a. Login to SSB as *root* locally (or remotely using SSH) to access the Console menu.

Step b. Select **Shells** > **Core Shell**, and issue the `syslog-ng stop` command.

Step c. Issue the `date` and check the system date and time. If it is incorrect (for example it displays 2000 January), replace the system battery. For details, see the hardware manual of the appliance.

Step d. Repeat the above steps on the other SSB node.

Step 2. *Optional step for data recovery*: Check the log spaces saved on the SSB nodes.

Step a. Login to the node from a local console.

Step b. Select **Shells** > **Core Shell** and enter `cd /opt/ssb/var/logspace/`. The log spaces are located under this directory.

Step c. Find which files were modified since the split brain situation occurred.

Step 3. Decide which node should be the master node from now on, then perform the following steps on the to-be-slave node:

Step a. Login to the node from a local console.

Step b. *Optional step for data recovery*: Backup the log messages that were modified since the split brain situation occurred.

Warning
This data will be deleted from the SSB node when the split-brain situation is resolved There is no way to import this data back into the database of SSB; it will be available only for offline use.

Step c. *Optional step for data recovery*: Type `exit` to return to the console menu.

Step d. Select **Shells** > **Boot shell**. If the to-be-slave node is not already the slave node, fail over the cluster to the other node manually by issuing the `/usr/share/heartbeat/hb_standby` command.

Step e. Stop the core firmware. Issue the `/etc/init.d/boot-xcb stop` command.

Step f. Invalidate the DRBD. Issue the following commands:

```
/sbin/drbdsetup /dev/drbd0 disconnect

/sbin/drbdsetup /dev/drbd0 invalidate.
```

Step 4. Reboot the to-be-slave node.

Step 5. Reboot the to-be-master node. The SSB cluster will be now functional, accepting traffic as before.

Step 6. After both nodes reboot, the cluster should be in **Degraded Sync** state, the master being **SyncSource** and the slave being **SyncTarget**. The master node should start synchronizing its data to the slave node. Depending on the amount of data, this can take a long time. To adjust the speed of the synchronization, see *Section 6.2.1, Adjusting the synchronization speed (p. 89)*.

16.6.4. Procedure – Replacing a node in an SSB HA cluster

Purpose:

To replace a unit in an SSB cluster with a new appliance, complete the following steps.

Steps:

Step 1. Verify the HA status on the working node. Select **Basic Settings** > **High Availability**. If one of the nodes has broken down or is missing, the **Status** field displays *DEGRADED*.

Step 2. Note down the IP addresses of the **Heartbeat** and the **Next hop monitoring** interfaces.

Step 3. Perform a full system backup. Before replacing the node, create a complete system backup of the working node. For details, see *Section 4.7, Data and configuration backups (p. 55)*.

Step 4. Check which firmware version is running on the working node. Select **Basic Settings** > **System** > **Version details** and write down the exact version numbers.

Step 5. Login to your MyBalaBit account at *https://my.balabit.com/login/* and download the CD ISO for the same SSB version that is running on your working node.

Step 6. Without connecting the replacement unit to the network, install the replacement unit from the ISO file. Use the IPMI interface if needed.

Step 7. When the installation is finished, connect the two SSB units with an Ethernet cable via the Ethernet connectors labeled as *4* (or *HA*).

Step 8. Reboot the replacement unit and wait until it finishes booting.

Step 9. Login to the working node and verify the HA state. Select **Basic Settings** > **High Availability**. The **Status** field should display *HALF*.

Step 10. Reconfigure the IP addresses of the **Heartbeat** and the **Next hop monitoring** interfaces. Click

.

Step 11. Click **Other node** > **Join HA**.

Step 12. Click **Other node > Reboot**.

Step 13. The replacement unit will reboot and start synchronizing data from the working node. The **Basic Settings > High Availability > Status** field will display *DEGRADED SYNC* until the synchronization finishes. Depending on the size of the hard disks and the amount of data stored, this can take several hours.

Step 14. After the synchronization is finished, connect the other Ethernet cables to their respective interfaces (external to *1* or *EXT*, management to *2* or *MGMT*) as needed for your environment.
Expected result:

A node of the SSB cluster is replaced with a new appliance.

16.6.5. Procedure – Resolving an IP conflict between cluster nodes

The IP addresses of the HA interfaces connecting the two nodes are detected automatically, during boot. When a node comes online, it attempts to connect to the IP address 1.2.4.1. If no other node responds until timeout, then it sets the IP address of its HA interface to 1.2.4.1, otherwise (if there is a responding node on 1.2.4.1) it sets its own HA interface to 1.2.4.2.

Replaced nodes do not yet know the HA configuration (or any other HA settings), and will attempt to negotiate it automatically in the same way. If the network is, for any reason, too slow to connect the nodes on time, the replacement node boots with the IP address of 1.2.4.1, which can cause an IP conflict if the other node has also set its IP to that same address previously. In this case, the replacement node cannot join the HA cluster.

To manually assign the correct IP address to the HA interface of a node, perform the following steps:

Step 1. Log in to the node using the IPMI interface or the physical console.
Configuration changes have not been synced to the new (replacement) node, as it could not join the HA cluster. Use the default password of the root user of SSB; see *Procedure B.1, Installing the SSB hardware (p. 238)*.

Step 2. From the console menu, choose **9 HA address**.

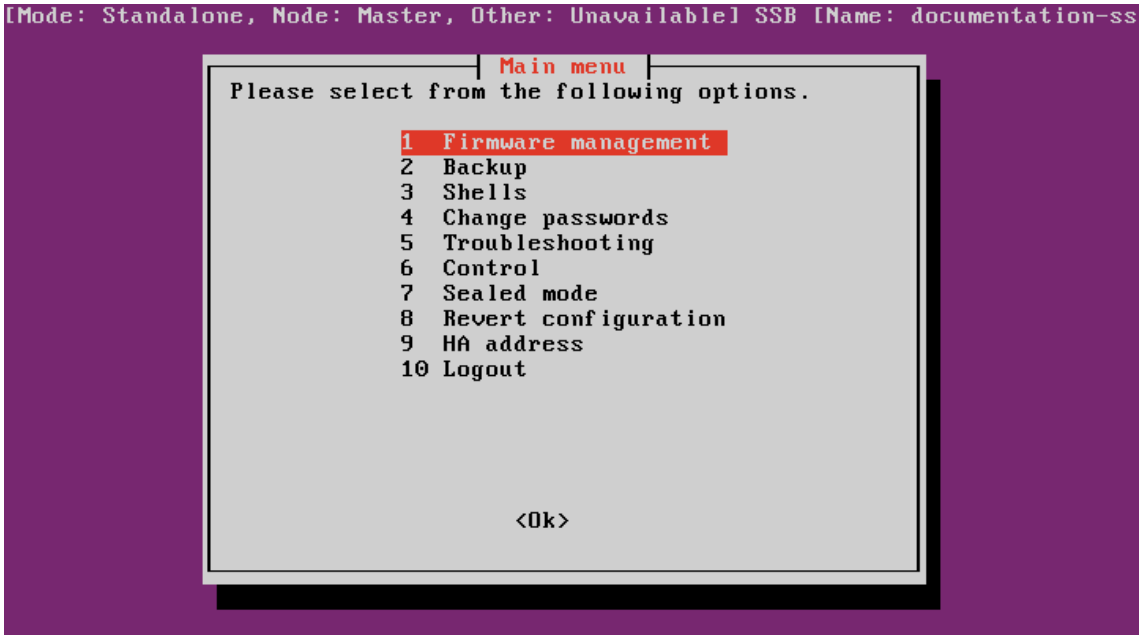

Figure 16.4. The console menu

Step 3. Choose the IP address of the node.

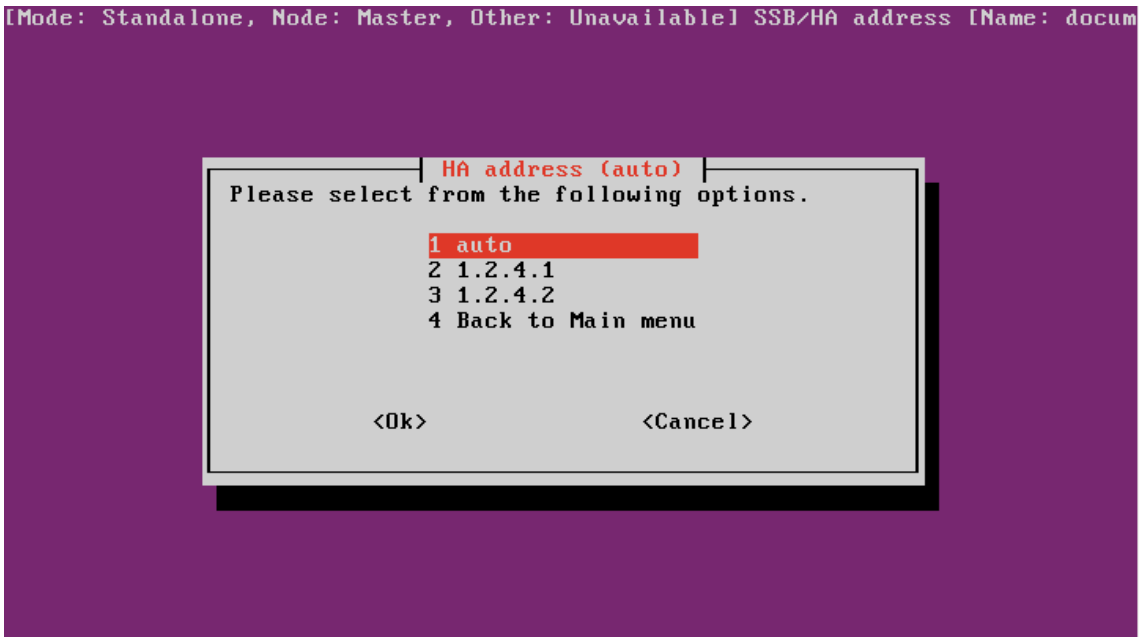
[Mode: Standalone, Node: Master, Other: Unavailable] SSB/HA address [Name: docum
HA address (auto)
Please select from the following options.
1 auto
2 1.2.4.1
3 1.2.4.2
4 Back to Main menu
<Ok>
<Cancel>

Figure 16.5. The console menu

Step 4. Reboot the node.

16.7. Procedure – Restoring SSB configuration and data

Purpose:

The following procedure describes how to restore the configuration and data of SSB from a complete backup, for example, after a hardware replacement.

Steps:

Step 1. Connect to your backup server and locate the directory where SSB saves the backups. The configuration backups are stored in the `config` subdirectory in timestamped files. Find the latest configuration file (the configuration files are called *`SSB-timestamp.config`*).

Step 2. Connect to SSB.

If you have not yet completed the Welcome Wizard, click **Browse**, select the configuration file, and click **Import**.

If you have already completed the Welcome Wizard, navigate to **Basic Settings** > **System** > **Import configuration** > **Browse**, select the configuration file, and click **Import**.

Step 3. Navigate to **Policies** > **Backup & Archive/Cleanup**. Verify that the settings of the target servers and the backup protocols are correct.

Step 4. Navigate to **Basic Settings** > **Management** > **System backup**, click **Restore now** and wait for the process to finish. Depending on the amount of data stored in the backup, and the speed of the connection to the backup server, this may take a long time.

Step 5. Navigate to **Log** > **Spaces**, and click **Restore ALL**. Depending on the amount of data stored in the backup, and the speed of the connection to the backup server, this may take a long time.

Appendix A. Package contents inventory

Carefully unpack all server components from the packing cartons. The following items should be packaged with the syslog-ng Store Box:

- A syslog-ng Store Box appliance, pre-installed with the latest syslog-ng Store Box firmware.
- syslog-ng Store Box accessory kit, including the following:
  - syslog-ng Store Box 4 F1 Packaging Checklist (this document).
  - GPL v2.0 license.
- Rack mount hardware.
- Power cable.

The default BIOS and IPMI passwords are in the documentation.

Appendix B. syslog-ng Store Box Hardware Installation Guide

This leaflet describes how to set up the syslog-ng Store Box (SSB) hardware. Refer to the following documents for step-by-step instructions:

- *syslog-ng Store Box T-1*: see the *SC512 Chassis Series User's Manual, Chapter 6: Rack Installation*. The manual is available online at *http://www.supermicro.com/manuals/chassis/1U/SC512.pdf*.
- *syslog-ng Store Box T-4*: see the *SC815 Chassis Series User's Manual, Chapter 6: Rack Installation*. The manual is available online at *http://www.supermicro.com/manuals/chassis/1U/SC815.pdf*.
- *syslog-ng Store Box T-10*: see the *SC219 Chassis Series User's Manual, Chapter 5: Rack Installation*. The manual is available online at *http://www.supermicro.com/manuals/chassis/2U/SC219.pdf*.

The manuals are also available online at *the BalaBit Documentation page*. Note that SSB hardware is built to custom specifications: CPU, memory, network card, and storage options differ from the stock chassis. You can find the hardware specifications in *Appendix C, Hardware specifications (p. 241)*

- For details on how to install a single SSB unit, see *Procedure B.1, Installing the SSB hardware (p. 238)*.
- For details on how to install a two SSB units in high availability mode, see *Procedure B.2, Installing two SSB units in HA mode (p. 240)*.

B.1. Procedure – Installing the SSB hardware

Purpose:

To install a single SSB unit, complete the following steps.

Steps:

Step 1. Unpack SSB.

Step 2. *Optional step*: Install SSB into a rack with the slide rails. Slide rails are available for all SSB appliances.

Step 3. Connect the cables.

Step a. Connect the Ethernet cable facing your LAN to the Ethernet connector labeled as `1`. This is the external interface of SSB and is used to configure SSB. (For details on the roles of the different interfaces, see *Section 2.6, Network interfaces (p. 8)*.)

Step b. Connect an Ethernet cable that you can use to remotely support the SSB hardware to the `IPMI` interface of SSB. For details, see the following documents:
The *Onboard BMC/IPMI User's Guide*, available on the BalaBit Hardware Documentation page at *https://www.balabit.com/support/documentation/*.

Warning
Connect the IPMI before plugging in the power chord. Failing to do so will result in IPMI failure.

It is not necessary for the IPMI interface to be accessible from the Internet, but the administrator of SSB must be able to access it for support and troubleshooting purposes in case vendor support is needed. The following ports are used by the IMPI interface:

- Port 623 (UDP): IPMI (cannot be changed)
- Port 5123 (UDP): floppy (cannot be changed)
- Port 5901 (TCP): video display (configurable)
- Port 5900 (TCP): HID (configurable)
- Port 5120 (TCP): CD (configurable)
- Port 80 (TCP): HTTP (configurable)

Access to information available only via the IPMI interface is a not mandatory, but highly recommended to speed up the support and troubleshooting processes.

Step c. *Optional step*: Connect the Ethernet cable to be used for managing SSB after its initial configuration to the Ethernet connector labeled as *2*. This is the management interface of SSB. (For details on the roles of the different interfaces, see *Section 2.6, Network interfaces (p. 8)*.)

Step d. *Optional step*: Connect the Ethernet cable connecting SSB to another SSB node to the Ethernet connector labeled as *4*. This is the high availability (HA) interface of SSB. (For details on the roles of the different interfaces, see *Section 2.6, Network interfaces (p. 8)*.)

Step 4. Power on the hardware.

Step 5. Change the BIOS and IPMI passwords on the syslog-ng Store Box. The default password is `ADMIN` or `changeme`, depending on your hardware.

Step 6. Following boot, SSB attempts to receive an IP address automatically via DHCP. If it fails to obtain an automatic IP address, it starts listening for HTTPS connections on the *192.168.1.1* IP address. To configure SSB to listen for connections on a custom IP address, complete the following steps:

Step a. Access SSB from the local console, and log in with username *root* and password *default*.

Step b. In the Console Menu, select **Shells** > **Core shell**.

Step c. Change the IP address of SSB:

```
ifconfig eth0 <IP-address> netmask 255.255.255.0
```

Replace `<IP-address>` with an IPv4 address suitable for your environment.

Step d. Set the default gateway using the following command:

```
route add default gw <IP-of-default-gateway>
```

Replace `<IP-of-default-gateway>` with the IP address of the default gateway.

Step e. Type `exit`, then select **Logout** from the Console Menu.

Step 7. Connect to the SSB web interface from a client machine and complete the Welcome Wizard as described in *Chapter 3, The Welcome Wizard and the first login (p. 16)*.

Note
The syslog-ng Store Box Administrator Guide is available on the SSB on the *BalaBit Documentation page*.

B.2. Procedure – Installing two SSB units in HA mode

Purpose:

To install SSB with high availability support, complete the following steps.

Steps:

Step 1. For the first SSB unit, complete *Procedure B.1, Installing the SSB hardware (p. 238)*.

Step 2. For the second SSB unit, complete Steps 1-3 of *Procedure B.1, Installing the SSB hardware (p. 238)*.

Step 3. Connect the two units with an Ethernet cable via the Ethernet connectors labeled as *4*.

Step 4. Power on the second unit.

Step 5. Change the BIOS and IPMI passwords on the second unit. The default password is `ADMIN` or `changeme`, depending on your hardware.

Step 6. Connect to the SSB web interface of the first unit from a client machine and enable the high availability mode. Navigate to **Basic Settings > High Availability** . Click **Convert to Cluster**, then reload the page in your browser.

Step 7. Click **Reboot Cluster**.

Step 8. Wait until the slave unit synchronizes its disk to the master unit. Depending on the size of the hard disks, this may take several hours. You can increase the speed of the synchronization via the SSB web interface at **Basic Settings > High Availability > DRBD sync rate limit**.

Appendix C. Hardware specifications

SSB appliances are built on high performance, energy efficient, and reliable hardware that are easily mounted into standard rack mounts.

Product	SSB T-1	SSB T-4	SSB T-10
Redundant PSU	No	Yes	Yes
Processor	Intel(R) Xeon(R) X3430 @ 2.40GHz	Intel(R) Xeon(R) E3-1275V2 @ 3.50GHz	2 x Intel(R) Xeon(R) E5-2630V2 @ 2.6GHz
Memory	2 x 4 GB	2 x 4 GB	8 x 4 GB
Capacity	2 x 1 TB	4 x 2 TB	13 x 1 TB
RAID	Software RAID	LSI MegaRAID SAS 9271-4i SGL	LSI 2208 (1GB cache)
IPMI	Yes	Yes	Yes
NIC	2x Intel® 82574L Gigabit Ethernet Controllers (Label 1, 2) Supermicro AOC-SG-i2 Dual GbE PCI-E x4 (Label 3, 4)	2x Intel® 82574L Gigabit Ethernet Controllers (Label 1, 2) Supermicro AOC-SG-i2 Dual GbE PCI-E x4 (Label 3, 4)	Intel® i350 Dual Port Gigabit Ethernet (Label 1, 2) Supermicro AOC-SG-i2 Dual GbE PCI-E x4 (Label 3, 4)

Table C.1. Hardware specifications

In addition to the network interface controllers listed above, the SSB T-10 appliance has two additional network interfaces (labels A, B). These interfaces will be used in a future release.

Appendix D. syslog-ng Store Box Software Installation Guide

This leaflet describes how to install the syslog-ng Store Box (SSB) software on a certified hardware. The list of certified hardware is available at BalaBit.

D.1. Procedure – Installing the SSB software

Purpose:

To install a new SSB on a server, complete the following steps:

Steps:

Step 1. Login to your *MyBalaBit account* and download the latest syslog-ng Store Box installation ISO file. Note that you need to have partner access to download syslog-ng Store Box ISO files. If you are a partner but do not see the ISO files, you can request partner access within MyBalaBit.

Step 2. Mount the ISO image, or burn it to a CD-ROM.

Step 3. Connect your computer to the `IPMI` interface of SSB. For details, see the following documents:
The *Onboard BMC/IPMI User's Guide*, available on the BalaBit Hardware Documentation page at *https://www.balabit.com/support/documentation/*.

Step 4. Power on the server.

Step 5. Login to the IPMI web interface, and boot the syslog-ng Store Box installation CD on the server using a virtual CD-ROM. For details, see the following documents:
The *Onboard BMC/IPMI User's Guide*, available on the BalaBit Hardware Documentation page at *https://www.balabit.com/support/documentation/*.

Step 6. When the syslog-ng Store Box installer starts, select **Installer**, press Enter, and wait until the server finishes the boot process.

Step 7. Select **Install a new SSB** and press **Enter** to start the installation process. Depending on the size of the disks, the installation process takes from a few minutes to an hour to complete. The progress of the installation is indicated in the **Installation Steps** window.

Step 8. The installer displays the following question: **Warning, all data on the hard drive(s) will be erased. Are you sure?** Select **Yes** and press **Enter**.

Step 9. The installer displays the MAC addresses of the network interfaces found in the SSB unit. Record these addresses.

Step 10. The installer displays the product name (the SSB configuration that was installed, for example, `SSB1000`). If the product name displayed does not match the product you wanted to install, complete the following steps:

Step a. Check that the hardware configuration of the appliance matches the specifications provided by BalaBit.

Step b. If the configuration matches the specifications but the installer displays a different product name, contact the BalaBit Support Team.

Step 11. During the **Finishing the Setup** step, the installer performs RAID synchronization.

- Select **Yes** to perform the RAID synchronization. RAID synchronization is a two-step process, the progress of the active step is indicated on the progress bar. Wait until both steps are completed. Note that this synchronization takes several hours (about 8 hours on the average).

- Select **No** to skip the RAID synchronization. Note that the system will automatically perform the synchronization after the first boot, but in this case the process will take several days.

Step 12. After the installation is finished, press Enter to return to the main menu.

Step 13. Select **Reboot** and press Enter to restart the system. Wait until the system reboots.

Step 14. Connect your computer to the *EXT* interface of SSB. Create an alias IP address for your computer that falls into the *192.168.1.0/24* subnet (for example *192.168.1.10*). For details, see *Section 3.1, The initial connection to SSB (p. 16)*.

Step 15. Open the *http://192.168.1.1* URL in your web browser and verify that the Welcome Wizard of the syslog-ng Store Box is available.

Note
For details on the supported web browsers and operating systems, see *Section 4.1, Supported web browsers and operating systems (p. 31)*.

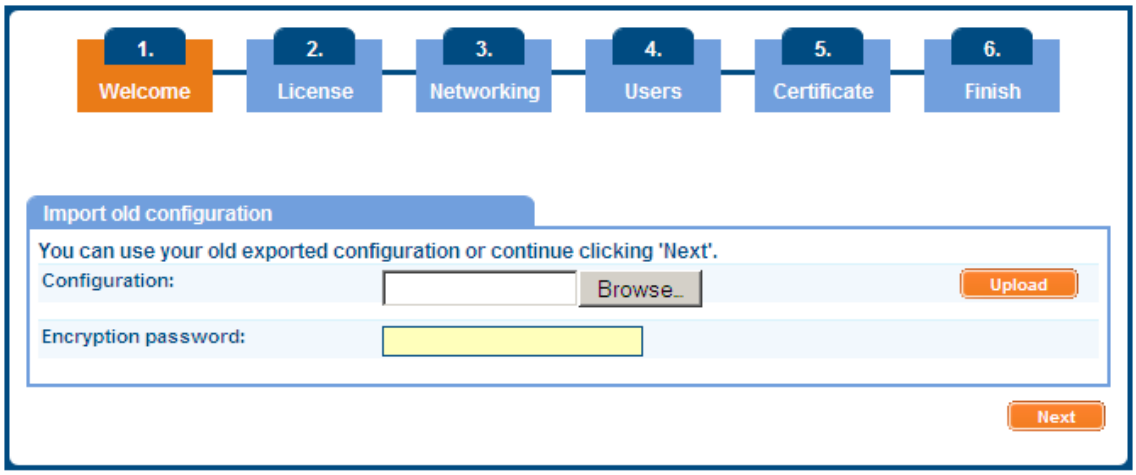

Figure D.1. The Welcome Wizard

Step 16. Power off the system.

Appendix E. syslog-ng Store Box VMware Installation Guide

This tutorial describes the possibilities and limitations of installing syslog-ng Store Box (SSB) 4 F1 as a virtual appliance under a VMware ESXi server.

E.1. Limitations of SSB under VMware

The following limitations apply to running version 4 F1 of SSB under VMware:

- SSB can be installed under the following VMware versions:
  - VMware ESXi 4.0 or later.
  - VMware ESX 4.0 or later.
  - VMware ESXi 5.0 or later.
- SSB can only use fixed disk space assigned to the virtual host; it is not possible to use on-demand disk allocation scenarios. To increase the size of the virtual disk, see *Procedure E.3, Modifying the virtual disk size under VMware (p. 245)*.
- High availability mode is not supported in VMware.
- Hardware-related alerts and status indicators of SSB may display inaccurate information, for example, display degraded RAID status.

E.2. Procedure – Installing SSB under VMware ESXi/ESX

Purpose:

To install a new SSB under VMware ESXi or ESX, complete the following steps:

Steps:

Step 1. Create the virtual machine for SSB using the following settings:

- Guest operating system: **Linux/Ubuntu 64-bit**.
- Allocate memory for the virtual machine. SSB requires a minimum of 1 GiB of memory, in addition to the memory limit of the indexed logspaces. The recommended size for the memory depends on the exact environment, but consider the following:
  - The base system requires 256 MiB.
  - The syslog-ng server running on SSB requires between 128 MiB and 1 GiB of memory, depending on the message load and on the configuration of SSB.
  - For every log space, SSB requires additional memory to index the incoming messages. The amount of memory allocated for the indexer can be set individually for every log space.
- The hard disk controller must be **LSI Logic Parallel**.

- Do not use RAID for the hard disk, use the data duplication features of your virtual environment instead. That way, a single hard disk is sufficient for the system. If you need to use the built-in RAID support of SSB for some reason, use two hard disks, and SSB will automatically use them in software RAID.

Warning
Hazard of data loss! When you install or reinstall SSB in a virtual environment, always create new hard disks. Using existing hard disks can cause unexpected behavior and operational problems.

- SSB uses a single fixed-sized disk:
  - About 5GB is required for the base system, the remaining disk space is used to store data. Use one of the following disk sizes (in GB): 8, 12, 68, 160, 230, 250, 465, 500, 925, 1029, 2053, 3077, 4101, 4645, 5125, 6149, 7173, 8197, 9205.
  - To increase the initial disk size, see *Procedure E.3, Modifying the virtual disk size under VMware (p. 245)*.
- SSB requires 4 network cards, all of them must be **E1000**.

Step 2. After creating the virtual machine, edit the settings of the machine. Set the following options:

Step a. Under **Options** > **VMware Tools** enable the **Shutdown, Suspend, Reset** options, otherwise the SSB administrator will not be able to access these functions from the SSB web interface.

Step b. Under **Options** > **Boot options** enable the **Force BIOS Setup** option. This is required to be able to check the system time (and modify it if needed) before installing SSB.

Step 3. Login to your *MyBalaBit account* and download the latest syslog-ng Store Box installation ISO file. Note that you need to have purchased SSB as a virtual appliance or have partner access to download syslog-ng Store Box ISO files. If you are a partner but do not see the ISO files, you can request partner access within MyBalaBit.

Step 4. Mount the ISO image and boot the virtual machine. Follow the on-screen instructions to install SSB.

E.3. Procedure – Modifying the virtual disk size under VMware

SSB can only use fixed disk space assigned to the virtual host. If you must increase the size of the virtual disk, complete the following steps:

Step 1. Create a full system backup (configuration and data backup). For detailed instructions, see *Section 4.7, Data and configuration backups (p. 55)*.

Step 2. Power down the virtual machine.

Step 3. Increase the storage size.

Step 4. Re-install SSB.

Step 5. Restore the system from the full backup. For detailed instructions, see *Procedure 16.7, Restoring SSB configuration and data (p. 235)*.

Appendix F. syslog-ng Store Box Hyper-V Installation Guide

This tutorial describes the possibilities and limitations of installing syslog-ng Store Box (SSB) 4 F1 as a virtual appliance under a Hyper-V server.

F.1. Limitations of SSB under Hyper-V

Version 4 F1 of SSB has no special support for running under Hyper-V. While the basic functionality of SSB is not affected by running as a virtual appliance, the following limitations apply:

- High availability mode is not supported in Hyper-V.
- Hardware-related alerts and status indicators of SSB may display inaccurate information, for example, display degraded RAID status.
- When running SSB under Microsoft Hyper-V, ensure that the network interfaces are actually connected to the network. When running under Hyper-V, SSB indicates on the **Basic Settings** > **Network** > **Ethernet links** page that there is a link even if the network interface is configured and enabled, but not connected to the network.
- When rebooting SSB in Hyper-V, the following critical error message may appear in the event log of the Hyper-V host:

  ```
  <Virtual machine name> was reset because an unrecoverable error occurred
  on a virtual processor that caused a triple fault.
  ```

  This is normal, there is no problem with SSB. For details, see *http://support.microsoft.com/kb/2711608*

F.2. Procedure – Installing SSB under Hyper-V

Purpose:

To install a new SSB under Hyper-V, complete the following steps:

Steps:

Step 1. Create the virtual machine for SSB using the following settings:

- Choose **Generation 1** for the virtual machine.
- Allocate memory for the virtual machine. SSB requires a minimum of 1 GiB of memory, in addition to the memory limit of the indexed logspaces. The recommended size for the memory depends on the exact environment, but consider the following:
  - The base system requires 256 MiB.
  - The syslog-ng server running on SSB requires between 128 MiB and 1 GiB of memory, depending on the message load and on the configuration of SSB.

- For every log space, SSB requires additional memory to index the incoming messages. The amount of memory allocated for the indexer can be set individually for every log space.

- Use one of the following disk sizes (in GB): 8, 12, 68, 160, 230, 250, 465, 500, 925, 1029, 2053, 3077, 4101, 4645, 5125, 6149, 7173, 8197, 9205.
  About 5GB is required for the base system, the remaining disk space is used to store data.

- Do not use RAID for the hard disk, use the data duplication features of your virtual environment instead. That way, a single hard disk is sufficient for the system. If you need to use the built-in RAID support of SSB for some reason, use two hard disks, and SSB will automatically use them in software RAID.

Warning
Hazard of data loss! When you install or reinstall SSB in a virtual environment, always create new hard disks. Using existing hard disks can cause unexpected behavior and operational problems.

- SSB requires 4 network cards. After completing the steps of the **New Virtual Machine Wizard**, add three additional network cards in the **Settings** of the virtual machine.

Step 2. Login to your *MyBalaBit account* and download the latest syslog-ng Store Box installation ISO file. Note that you need to have purchased SSB as a virtual appliance or have partner access to download syslog-ng Store Box ISO files. If you are a partner but do not see the ISO files, you can request partner access within MyBalaBit.

Step 3. Mount the ISO image and boot the virtual machine. Follow the on-screen instructions to install SSB.

Appendix G. License contract for BalaBit Product

SUBJECT OF THE LICENSE CONTRACT

This License Contract is entered into by and between BalaBit S.a.r.l. (or, based on your place of operation, one of its affiliates) as Licensor (hereinafter Company or Licensor, or BalaBit) and Licensee (hereinafter Licensee) and sets out the terms and conditions under which Licensee and/or Licensee's Authorized Subsidiaries may use the BalaBit product under this License Contract.

DEFINITIONS

In this License Contract, the following words shall have the following meanings:

Name	Description
Annexed Software	Any third party software that is a not a BalaBit Product contained in the install media of the BalaBit Product.
BalaBit Product	Any software, hardware or service Licensed, sold, or provided by BalaBit including any installation, education, support and warranty services or any product falling under the copyright of BalaBit with the exception of the Annexed Software.
License Contract	The present BalaBit Product License Contract.
Product Documentation	Any documentation referring to the BalaBit Product or any module thereof, with special regard to the administration guide, the product description, the installation guide, user guides and manuals.
End-user Certificate	The document signed by Licensor which contains a) identification data of Licensee; b) configuration of BalaBit Product and designation of Licensed modules thereof; c) declaration of the parties on accepting the terms and conditions of this License Contract; and d) declaration of Licensee that is in receipt of the install media and the hardware appliance.
Product Usage Terms	Defines the conditions (related usage environment and limitations) under the BalaBit Product may be used by the Licensee.
Warranty Period	The period of twelve (12) months from the date of delivery of the BalaBit Product to Licensee.

Table G.1. Words and expressions

LICENSE GRANTS AND RESTRICTIONS

Based on the terms and conditions of the present License Contract and the applicable End-user Certificate, and the Product Usage Terms BalaBit grants to Licensee, a non-exclusive, perpetual license to use BalaBit Product. License is transferable only with the prior written approval of BalaBit.

Licensee shall use the BalaBit Product in accordance with the conditions sets by the Product Usage Terms, especially in the configuration and in the quantities specified in the End-user Certificate and Product Usage Terms.

On the install media (firmware CD-ROM, USB stick) all modules of the BalaBit Product will be presented, however, Licensee shall not be entitled to use any module which was not Licensed to it. Access rights to modules and IP connections are controlled by an "electronic key" accompanying the BalaBit Product.

Licensee shall be entitled to make one back-up copy of the install media containing the BalaBit Product.

Licensee shall make the BalaBit Product available solely to its own employees and those of the Authorized Subsidiaries.

Licensee shall take all reasonable steps to protect BalaBit's rights with respect to the BalaBit Product with special regard and care to protecting it from any unauthorized access.

Licensee shall, in 5 working days, properly answer the queries of BalaBit referring to the actual usage conditions of the BalaBit Product that may differ or allegedly differs from the License conditions.

Licensee shall not modify the BalaBit Product in any way, with special regard to the functions that inspect the usage of the software. Licensee shall install the code permitting the usage of the BalaBit Product according to the provisions defined for it by BalaBit. Licensee may not modify or cancel such codes. Configuration settings of the BalaBit Product in accordance with the possibilities offered by the system shall not be construed as modification of the software.

Licensee shall only be entitled to analyze the structure of the BalaBit Products (decompilation or reverse-engineering) if concurrent operation with a software developed by a third party is necessary, and upon request to supply the information required for concurrent operation BalaBit does not provide such information within 60 days from the receipt of such a request.

These user actions are limited to parts of the BalaBit Product which are necessary for concurrent operation. Any information obtained as a result of applying the previous Section (i) cannot be used for purposes other than concurrent operation with the BalaBit Product; (ii) cannot be disclosed to third parties unless it is necessary for concurrent operation with the BalaBit Product; (iii) cannot be used for the development, production or distribution of a different software which is similar to the BalaBit Product in its form of expression, or for any other act violating copyright.

For any Annexed Software contained by the same install media as the BalaBit Product, the terms and conditions defined by its copyright owner shall be properly applied. BalaBit does not grant any License rights to any Annexed Software.

Any usage of the BalaBit Product exceeding the limits and restrictions defined in this License Contract shall qualify as material breach of the License Contract and Licensee shall be fully liable towards BalaBit.

Licensee shall have the right to obtain and use content updates only if Licensee concludes a support contract that includes such content updates (maintenance of the software), or if Licensee has otherwise separately acquired the right to obtain and use such content updates. This License Contract does not otherwise permit Licensee to obtain and use content updates.

INTELLECTUAL PROPERTY RIGHTS

Licensee agrees that BalaBit owns all rights, titles, and interests related to the BalaBit Product, including all of BalaBit's patents, trademarks, trade names, inventions, economic intellectual property rights, know-how, and trade secrets relating to the design, manufacture, operation or service of the BalaBit Products.

The use by Licensee of any of these intellectual property rights is authorized only for the purposes set forth herein, and upon termination of this License Contract for any reason, such authorization shall cease.

The BalaBit Products are licensed only for the Licensee’s own internal business purposes in every case, under the condition that such License does not convey any license, expressly or by implication, to manufacture duplicate or otherwise copy or reproduce any of the BalaBit Products. The sublicense to third parties, or provision of any services with the utilization of the BalaBit Product is not allowed. No other rights than expressly stated herein are granted to Licensee.

Licensee will take appropriate steps with its Authorized Subsidiaries, as BalaBit may request, to inform them of and assure compliance with the restrictions contained in the License Contract.

WARRANTIES

BalaBit warrants that during the Warranty Period, the magnetic or optical media upon which the BalaBit Product is recorded will not be defective under normal use. BalaBit will replace any defective media returned to it, accompanied by a dated proof of purchase, within the Warranty Period at no charge to Licensee. Upon receipt of the allegedly defective BalaBit Product, BalaBit will at its option, deliver a replacement BalaBit Product or BalaBit's current equivalent Product to Licensee at no additional cost. BalaBit will bear the delivery charges to Licensee for the replacement Product.

Should BalaBit Product be used in conjunction with third party software, BalaBit shall not be liable for errors due to third party software.

BalaBit warrants, that during the Warranty Period, the BalaBit Product, under normal use in the operating environment defined by BalaBit, and without unauthorized modification, will perform in substantial compliance with the Product Documentation accompanying the BalaBit Product, when used on that hardware for which it was installed, in compliance with the provisions of the user manuals and the recommendations of BalaBit.

The date of the notification sent to BalaBit shall qualify as the date of the failure. Licensee shall do its best to mitigate the consequences of that failure. If, during the Warranty Period, the BalaBit Product fails to comply with this warranty, and such failure is reported by Licensee to BalaBit within the Warranty Period, BalaBit's sole obligation and liability for breach of this warranty is, at BalaBit's sole option, either: (i) to correct such failure, (ii) to replace the defective BalaBit Product or (iii) to refund the license fees paid by Licensee for the applicable BalaBit Product.

BREACH OF CONTRACT

In case of breach of contract with respect to BalaBit, or the BalaBit Product, committed by violating any provision of the present License Contract, Licensee shall pay liquidated damages to BalaBit. The amount of the liquidated damages shall be twice as much as the price of the BalaBit Product concerned, on BalaBit's current Price List.

INTELLECTUAL PROPERTY INDEMNIFICATION

BalaBit shall pay all damages, costs and reasonable attorney's fees awarded against Licensee in connection with any claim brought against Licensee to the extent that such claim is based on a claim that Licensee's authorized use of the BalaBit Product infringes a patent, copyright, trademark or trade secret. Licensee shall notify BalaBit in writing of any such claim as soon as Licensee learns of it and shall cooperate fully with BalaBit in connection with the defense of that claim. BalaBit shall have sole control of that defense (including without limitation the right to settle the claim).

If Licensee is prohibited from using any BalaBit Product due to an infringement claim, or if BalaBit believes that any BalaBit Product is likely to become the subject of an infringement claim, BalaBit shall at its sole option, either: (i) obtain the right for Licensee to continue to use such BalaBit Product, (ii) replace or modify the BalaBit Product so as to make such BalaBit Product non-infringing and substantially comparable in functionality or (iii) refund to Licensee the amount paid for such infringing BalaBit Product and provide a pro-rated refund of any unused, prepaid maintenance fees paid by Licensee, in exchange for Licensee's return of such BalaBit Product to BalaBit.

Notwithstanding the above, BalaBit will have no liability for any infringement claim to the extent that it is based upon: (i) modification of the BalaBit Product other than by BalaBit, (ii) use of the BalaBit Product in combination with any product not specifically authorized by BalaBit to be combined with the BalaBit Product or (iii) use of the BalaBit Product in an unauthorized manner for which it was not designed.

LICENSE FEE

The End-user Certificate and the Product Usage Term contain the details of the purchased License and usage limitations. This information serves as the calculation base of the License fee.

Licensee acknowledges that payment of the License fees is a condition of lawful usage.

License fees do not contain any installation or post charges, taxes, duties, etc.

The license right of BalaBit Product is transferred to the Licensee only when Licensee pays the License fee to BalaBit. In case of non-payment BalaBit has right to terminate, or rescind from the License Contract with immediate effect and Licensee has to send back the BalaBit Product on their own cost and takes all liability regarding the unlawful usage and the early termination.

DISCLAIMER OF WARRANTIES

EXCEPT AS SET OUT IN THIS LICENSE CONTRACT, BALABIT MAKES NO WARRANTIES OF ANY KIND WITH RESPECT TO THE BALABIT PRODUCT. TO THE MAXIMUM EXTENT PERMITTED BY APPLICABLE LAW, BALABIT EXCLUDES ANY OTHER WARRANTIES, INCLUDING BUT NOT LIMITED TO ANY IMPLIED WARRANTIES OF SATISFACTORY QUALITY, MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NON-INFRINGEMENT OF INTELLECTUAL PROPERTY RIGHTS.

LIMITATION OF LIABILITY

SOME STATES AND COUNTRIES, INCLUDING MEMBER COUNTRIES OF THE EUROPEAN UNION, DO NOT ALLOW THE LIMITATION OR EXCLUSION OF LIABILITY FOR INCIDENTAL OR CONSEQUENTIAL DAMAGES AND, THEREFORE, THE FOLLOWING LIMITATION OR EXCLUSION MAY NOT APPLY TO THIS LICENSE CONTRACT IN THOSE STATES AND COUNTRIES. TO THE MAXIMUM EXTENT PERMITTED BY APPLICABLE LAW AND REGARDLESS OF WHETHER ANY

REMEDY SET OUT IN THIS LICENSE CONTRACT FAILS OF ITS ESSENTIAL PURPOSE, IN NO EVENT SHALL BALABIT BE LIABLE TO LICENSEE FOR ANY SPECIAL, CONSEQUENTIAL, INDIRECT OR SIMILAR DAMAGES OR LOST PROFITS OR LOST DATA ARISING OUT OF THE USE OR INABILITY TO USE THE BALABIT PRODUCT EVEN IF BALABIT HAS BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES.

IN NO CASE SHALL BALABIT'S TOTAL LIABILITY UNDER THIS LICENSE CONTRACT EXCEED THE FEES PAID BY LICENSEE FOR THE BALABIT PRODUCT LICENSED UNDER THIS LICENSE CONTRACT.

BalaBit shall not take any responsibility, for damages caused by the usage of the BalaBit Product which is not in accordance with the Product Usage Terms.

DURATION AND TERMINATION

This License Contract shall come into effect on the day when the End-user Certificate and the declaration of the Licensee on accepting the terms and conditions of this License Contract and the declaration of Licensee that is in receipt of the install media and the hardware appliance are both signed by the duly authorized representatives of the relevant party.

Licensee may terminate the License Contract at any time by written notice sent to BalaBit and by simultaneously destroying all copies of the BalaBit Product licensed under this License Contract.

BalaBit may terminate this License Contract with immediate effect by written notice to Licensee, if Licensee is in material or persistent breach of the License Contract and either that breach is incapable of remedy or Licensee shall have failed to remedy that breach within 30 days after receiving written notice requiring it to remedy that breach.

AMENDMENTS

Save as expressly provided in this License Contract, no amendment or variation of this License Contract shall be effective unless in writing and signed by a duly authorized representative of the parties to it.

WAIVER

The failure of a party to exercise or enforce any right under this License Contract shall not be deemed to be a waiver of that right nor operate to bar the exercise or enforcement of it at any time or times thereafter.

SEVERABILITY

If any part of this License Contract becomes invalid, illegal or unenforceable, the parties shall in such an event negotiate in good faith in order to agree on the terms of a mutually satisfactory provision to be substituted for the invalid, illegal or unenforceable provision which as nearly as possible validly gives effect to their intentions as expressed in this License Contract.

NOTICES

Any notice required to be given pursuant to this License Contract shall be in writing and shall be given by delivering the notice by hand, or by sending the same by prepaid first class post (airmail if to an address outside the country of posting) or via e-mail to the address of the relevant party. Any notice given according to the above procedure shall be deemed to have been given at the time of delivery (if delivered by hand) and when received (if sent by post or via e-mail).

APPLICABLE LAW AND LANGUAGE AND SETTLEMENT OF LEGAL DISPUTES

This agreement shall be governed by and construed in accordance with the laws of the Hungary without regard for its conflicts of law provisions. The language for all communications regarding this Agreement shall be English. The Parties agree that any dispute arising from this Agreement, or the breach, termination, validity or interpretation thereof or in relation thereto shall come under the exclusive jurisdiction of the Hungarian Court as defined below.

The Parties irrevocably agree that any dispute, controversy or claim arising out of or in connection with this agreement, or the breach, termination or invalidity thereof, shall be finally settled by arbitration by the Permanent Arbitration Court attached to the Hungarian Chamber of Commerce and Industry acting in accordance with its own Rules of Procedure. The place of arbitration shall be Budapest, the number of arbitrators shall be three (3) and the language to be used in the arbitral proceedings shall be English.

The Arbitration Court shall consist of three arbitrators, out of which one shall act as chairman. The chairman should have the competence of judgeship. The Arbitration Court shall be created in the manner that the prosecuting Party, indicating the subject of the debate and nominating an arbitrator, calls the counterparty in writing to nominate the other arbitrator and the nominated arbitrators will elect the chairman.

MISCELLANEOUS

Headings are for convenience only and shall be ignored in interpreting this License Contract.

This License Contract and the rights granted in this License Contract may not be assigned, sublicensed or otherwise transferred in whole or in part by Licensee without BalaBit's prior written consent.

An independent third party auditor, reasonably acceptable to BalaBit and Licensee, may upon reasonable notice to Licensee and during normal business hours, but not more often than once each year, inspect Licensee's relevant records in order to confirm that usage of the BalaBit Product complies with the terms and conditions of this License Contract. BalaBit shall bear the costs of such audit. All audits shall be subject to the reasonable safety and security policies and procedures of Licensee.

In case of non-acceptance in the person, an auditor, appointed by BalaBit shall keep full, complete and accurate inspect concerning that usage of the BalaBit Product complies with the terms and conditions of this License Contract.

The auditor shall be entitled to examine, inspect, copy and audit the usage of the BalaBit Product. If the inspection or audit reveals that the usage is not complies with the conditions of the License Contract the Licensee shall immediately:

- (a) pay to BalaBit the amount of any underpayment, together with interest on that amount calculated at the rate of two per cent (2%) over the Barclay Bank base rate from time to time; and
- (b) pay the costs of the audit and/or inspection where that audit or inspection reveals an underpayment in excess of five per cent (5%).

In case of the License shall not let the auditor to inspect, or examine the usage of BalaBit Product, BalaBit has right to terminate or rescind from the License Contract with immediate effect and Licensee has to send back the BalaBit Product on their own cost and takes all liability regarding the unlawful usage and the early termination.

This License Contract constitutes the entire agreement between the parties with regard to the subject matter hereof.

Appendix H. Creative Commons Attribution Non-commercial No Derivatives (by-nc-nd) License

THE WORK (AS DEFINED BELOW) IS PROVIDED UNDER THE TERMS OF THIS CREATIVE COMMONS PUBLIC LICENSE ("CCPL" OR "LICENSE"). THE WORK IS PROTECTED BY COPYRIGHT AND/OR OTHER APPLICABLE LAW. ANY USE OF THE WORK OTHER THAN AS AUTHORIZED UNDER THIS LICENSE OR COPYRIGHT LAW IS PROHIBITED. BY EXERCISING ANY RIGHTS TO THE WORK PROVIDED HERE, YOU ACCEPT AND AGREE TO BE BOUND BY THE TERMS OF THIS LICENSE. TO THE EXTENT THIS LICENSE MAY BE CONSIDERED TO BE A CONTRACT, THE LICENSOR GRANTS YOU THE RIGHTS CONTAINED HERE IN CONSIDERATION OF YOUR ACCEPTANCE OF SUCH TERMS AND CONDITIONS.

1. *Definitions*
    a. "Adaptation" means a work based upon the Work, or upon the Work and other pre-existing works, such as a translation, adaptation, derivative work, arrangement of music or other alterations of a literary or artistic work, or phonogram or performance and includes cinematographic adaptations or any other form in which the Work may be recast, transformed, or adapted including in any form recognizably derived from the original, except that a work that constitutes a Collection will not be considered an Adaptation for the purpose of this License. For the avoidance of doubt, where the Work is a musical work, performance or phonogram, the synchronization of the Work in timed-relation with a moving image ("synching") will be considered an Adaptation for the purpose of this License.
    b. "Collection" means a collection of literary or artistic works, such as encyclopedias and anthologies, or performances, phonograms or broadcasts, or other works or subject matter other than works listed in Section 1(f) below, which, by reason of the selection and arrangement of their contents, constitute intellectual creations, in which the Work is included in its entirety in unmodified form along with one or more other contributions, each constituting separate and independent works in themselves, which together are assembled into a collective whole. A work that constitutes a Collection will not be considered an Adaptation (as defined above) for the purposes of this License.
    c. "Distribute" means to make available to the public the original and copies of the Work through sale or other transfer of ownership.
    d. "Licensor" means the individual, individuals, entity or entities that offer(s) the Work under the terms of this License.
    e. "Original Author" means, in the case of a literary or artistic work, the individual, individuals, entity or entities who created the Work or if no individual or entity can be identified, the publisher; and in addition (i) in the case of a performance the actors, singers, musicians, dancers, and other persons who act, sing, deliver, declaim, play in, interpret or otherwise perform literary or artistic works or expressions of folklore; (ii) in the case of a phonogram the producer being the person or legal entity who first fixes the sounds of a performance or other sounds; and, (iii) in the case of broadcasts, the organization that transmits the broadcast.

f. "Work" means the literary and/or artistic work offered under the terms of this License including without limitation any production in the literary, scientific and artistic domain, whatever may be the mode or form of its expression including digital form, such as a book, pamphlet and other writing; a lecture, address, sermon or other work of the same nature; a dramatic or dramatico-musical work; a choreographic work or entertainment in dumb show; a musical composition with or without words; a cinematographic work to which are assimilated works expressed by a process analogous to cinematography; a work of drawing, painting, architecture, sculpture, engraving or lithography; a photographic work to which are assimilated works expressed by a process analogous to photography; a work of applied art; an illustration, map, plan, sketch or three-dimensional work relative to geography, topography, architecture or science; a performance; a broadcast; a phonogram; a compilation of data to the extent it is protected as a copyrightable work; or a work performed by a variety or circus performer to the extent it is not otherwise considered a literary or artistic work.

g. "You" means an individual or entity exercising rights under this License who has not previously violated the terms of this License with respect to the Work, or who has received express permission from the Licensor to exercise rights under this License despite a previous violation.

h. "Publicly Perform" means to perform public recitations of the Work and to communicate to the public those public recitations, by any means or process, including by wire or wireless means or public digital performances; to make available to the public Works in such a way that members of the public may access these Works from a place and at a place individually chosen by them; to perform the Work to the public by any means or process and the communication to the public of the performances of the Work, including by public digital performance; to broadcast and rebroadcast the Work by any means including signs, sounds or images.

i. "Reproduce" means to make copies of the Work by any means including without limitation by sound or visual recordings and the right of fixation and reproducing fixations of the Work, including storage of a protected performance or phonogram in digital form or other electronic medium.

2. *Fair Dealing Rights.* Nothing in this License is intended to reduce, limit, or restrict any uses free from copyright or rights arising from limitations or exceptions that are provided for in connection with the copyright protection under copyright law or other applicable laws.

3. *License Grant.* Subject to the terms and conditions of this License, Licensor hereby grants You a worldwide, royalty-free, non-exclusive, perpetual (for the duration of the applicable copyright) license to exercise the rights in the Work as stated below:

   a. to Reproduce the Work, to incorporate the Work into one or more Collections, and to Reproduce the Work as incorporated in the Collections; and,

   b. to Distribute and Publicly Perform the Work including as incorporated in Collections.

   The above rights may be exercised in all media and formats whether now known or hereafter devised. The above rights include the right to make such modifications as are technically necessary to exercise the rights in other media and formats, but otherwise you have no rights to make Adaptations. Subject to 8(f), all rights not expressly granted by Licensor are hereby reserved, including but not limited to the rights set forth in Section 4(d).

4. *Restrictions.* The license granted in Section 3 above is expressly made subject to and limited by the following restrictions:

a. You may Distribute or Publicly Perform the Work only under the terms of this License. You must include a copy of, or the Uniform Resource Identifier (URI) for, this License with every copy of the Work You Distribute or Publicly Perform. You may not offer or impose any terms on the Work that restrict the terms of this License or the ability of the recipient of the Work to exercise the rights granted to that recipient under the terms of the License. You may not sublicense the Work. You must keep intact all notices that refer to this License and to the disclaimer of warranties with every copy of the Work You Distribute or Publicly Perform. When You Distribute or Publicly Perform the Work, You may not impose any effective technological measures on the Work that restrict the ability of a recipient of the Work from You to exercise the rights granted to that recipient under the terms of the License. This Section 4(a) applies to the Work as incorporated in a Collection, but this does not require the Collection apart from the Work itself to be made subject to the terms of this License. If You create a Collection, upon notice from any Licensor You must, to the extent practicable, remove from the Collection any credit as required by Section 4(c), as requested.

b. You may not exercise any of the rights granted to You in Section 3 above in any manner that is primarily intended for or directed toward commercial advantage or private monetary compensation. The exchange of the Work for other copyrighted works by means of digital file-sharing or otherwise shall not be considered to be intended for or directed toward commercial advantage or private monetary compensation, provided there is no payment of any monetary compensation in connection with the exchange of copyrighted works.

c. If You Distribute, or Publicly Perform the Work or Collections, You must, unless a request has been made pursuant to Section 4(a), keep intact all copyright notices for the Work and provide, reasonable to the medium or means You are utilizing: (i) the name of the Original Author (or pseudonym, if applicable) if supplied, and/or if the Original Author and/or Licensor designate another party or parties (for example a sponsor institute, publishing entity, journal) for attribution ("Attribution Parties") in Licensor's copyright notice, terms of service or by other reasonable means, the name of such party or parties; (ii) the title of the Work if supplied; (iii) to the extent reasonably practicable, the URI, if any, that Licensor specifies to be associated with the Work, unless such URI does not refer to the copyright notice or licensing information for the Work. The credit required by this Section 4(c) may be implemented in any reasonable manner; provided, however, that in the case of a Collection, at a minimum such credit will appear, if a credit for all contributing authors of Collection appears, then as part of these credits and in a manner at least as prominent as the credits for the other contributing authors. For the avoidance of doubt, You may only use the credit required by this Section for the purpose of attribution in the manner set out above and, by exercising Your rights under this License, You may not implicitly or explicitly assert or imply any connection with, sponsorship or endorsement by the Original Author, Licensor and/or Attribution Parties, as appropriate, of You or Your use of the Work, without the separate, express prior written permission of the Original Author, Licensor and/or Attribution Parties.

d. For the avoidance of doubt:

   i. Non-waivable Compulsory License Schemes. In those jurisdictions in which the right to collect royalties through any statutory or compulsory licensing scheme cannot be waived, the Licensor reserves the exclusive right to collect such royalties for any exercise by You of the rights granted under this License;

   ii. Waivable Compulsory License Schemes. In those jurisdictions in which the right to collect royalties through any statutory or compulsory licensing scheme can be waived, the Licensor reserves the exclusive right to collect such royalties for any exercise by You of the rights

granted under this License if Your exercise of such rights is for a purpose or use which is otherwise than noncommercial as permitted under Section 4(b) and otherwise waives the right to collect royalties through any statutory or compulsory licensing scheme; and,

iii. Voluntary License Schemes. The Licensor reserves the right to collect royalties, whether individually or, in the event that the Licensor is a member of a collecting society that administers voluntary licensing schemes, via that society, from any exercise by You of the rights granted under this License that is for a purpose or use which is otherwise than noncommercial as permitted under Section 4(b).

e. Except as otherwise agreed in writing by the Licensor or as may be otherwise permitted by applicable law, if You Reproduce, Distribute or Publicly Perform the Work either by itself or as part of any Collections, You must not distort, mutilate, modify or take other derogatory action in relation to the Work which would be prejudicial to the Original Author's honor or reputation.

5. *Representations, Warranties and Disclaimer* UNLESS OTHERWISE MUTUALLY AGREED BY THE PARTIES IN WRITING, LICENSOR OFFERS THE WORK AS-IS AND MAKES NO REPRESENTATIONS OR WARRANTIES OF ANY KIND CONCERNING THE WORK, EXPRESS, IMPLIED, STATUTORY OR OTHERWISE, INCLUDING, WITHOUT LIMITATION, WARRANTIES OF TITLE, MERCHANTIBILITY, FITNESS FOR A PARTICULAR PURPOSE, NONINFRINGEMENT, OR THE ABSENCE OF LATENT OR OTHER DEFECTS, ACCURACY, OR THE PRESENCE OF ABSENCE OF ERRORS, WHETHER OR NOT DISCOVERABLE. SOME JURISDICTIONS DO NOT ALLOW THE EXCLUSION OF IMPLIED WARRANTIES, SO SUCH EXCLUSION MAY NOT APPLY TO YOU.

6. *Limitation on Liability.* EXCEPT TO THE EXTENT REQUIRED BY APPLICABLE LAW, IN NO EVENT WILL LICENSOR BE LIABLE TO YOU ON ANY LEGAL THEORY FOR ANY SPECIAL, INCIDENTAL, CONSEQUENTIAL, PUNITIVE OR EXEMPLARY DAMAGES ARISING OUT OF THIS LICENSE OR THE USE OF THE WORK, EVEN IF LICENSOR HAS BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES.

7. *Termination*

   a. This License and the rights granted hereunder will terminate automatically upon any breach by You of the terms of this License. Individuals or entities who have received Collections from You under this License, however, will not have their licenses terminated provided such individuals or entities remain in full compliance with those licenses. Sections 1, 2, 5, 6, 7, and 8 will survive any termination of this License.

   b. Subject to the above terms and conditions, the license granted here is perpetual (for the duration of the applicable copyright in the Work). Notwithstanding the above, Licensor reserves the right to release the Work under different license terms or to stop distributing the Work at any time; provided, however that any such election will not serve to withdraw this License (or any other license that has been, or is required to be, granted under the terms of this License), and this License will continue in full force and effect unless terminated as stated above.

8. *Miscellaneous*

   a. Each time You Distribute or Publicly Perform the Work or a Collection, the Licensor offers to the recipient a license to the Work on the same terms and conditions as the license granted to You under this License.

   b. If any provision of this License is invalid or unenforceable under applicable law, it shall not affect the validity or enforceability of the remainder of the terms of this License, and without further

action by the parties to this agreement, such provision shall be reformed to the minimum extent necessary to make such provision valid and enforceable.

c. No term or provision of this License shall be deemed waived and no breach consented to unless such waiver or consent shall be in writing and signed by the party to be charged with such waiver or consent.

d. This License constitutes the entire agreement between the parties with respect to the Work licensed here. There are no understandings, agreements or representations with respect to the Work not specified here. Licensor shall not be bound by any additional provisions that may appear in any communication from You. This License may not be modified without the mutual written agreement of the Licensor and You.

e. The rights granted under, and the subject matter referenced, in this License were drafted utilizing the terminology of the Berne Convention for the Protection of Literary and Artistic Works (as amended on September 28, 1979), the Rome Convention of 1961, the WIPO Copyright Treaty of 1996, the WIPO Performances and Phonograms Treaty of 1996 and the Universal Copyright Convention (as revised on July 24, 1971). These rights and subject matter take effect in the relevant jurisdiction in which the License terms are sought to be enforced according to the corresponding provisions of the implementation of those treaty provisions in the applicable national law. If the standard suite of rights granted under applicable copyright law includes additional rights not granted under this License, such additional rights are deemed to be included in the License; this License is not intended to restrict the license of any rights under applicable law.

Glossary

alias IP	An additional IP address assigned to an interface that already has an IP address. The normal and alias IP addresses both refer to the same physical interface.
auditing policy	The auditing policy determines which events are logged on host running Microsoft Windows operating systems.
authentication	The process of verifying the authenticity of a user or client before allowing access to a network system or service.
BSD-syslog protocol	The old syslog protocol standard described in RFC 3164 *The BSD syslog Protocol*. Sometimes also referred to as the legacy-syslog protocol.
CA	A Certificate Authority (CA) is an institute that issues certificates.
certificate	A certificate is a file that uniquely identifies its owner. Certificates contains information identifying the owner of the certificate, a public key itself, the expiration date of the certificate, the name of the CA that signed the certificate, and some other data.
client mode	In client mode, syslog-ng collects the local logs generated by the host and forwards them through a network connection to the central syslog-ng server or to a relay.
destination	A logspace or a remote database or server where the log messages are stored.
destination driver	A communication method that syslog-ng uses to send log messages to a destination, for example to a remote server or to the hard disk.
destination, remote	A destination that sends log messages to a remote host (that is, a syslog-ng relay or server) using a network connection.
destination, local	A destination that transfers log messages to a logspace.
disk buffer	The Premium Edition of syslog-ng can store messages on the local hard disk if the central log server or the network connection to the server becomes unavailable.
disk queue	See *disk buffer*.
domain name	The name of a network, for example `balabit.com`.
External network interface	The *external* interface (labeled `1` or `EXT`) is used for general communication between the clients and the servers. If the management interface is not configured, the external interface is used for management purposes as well.

filter	An expression that selects only those message from a source that match the conditions set in the filter.
firmware	A firmware is a collection of the software components running on SSB. Individual software components cannot be upgraded on SSB, only the entire firmware. SSB contains two firmwares, an external (or boot) firmware and an internal (or core) firmware. These can be upgraded separately.
gateway	A device that connect two or more parts of the network, for example your local intranet and the external network (the Internet). Gateways act as entrances into other networks.
High Availability	High Availability (HA) uses a second SSB unit (called slave node) to ensure that the services are available even if the first unit (called master node) breaks down.
host	A computer connected to the network.
hostname	A name that identifies a host on the network. Hostnames can contain only alphanumerical characters (A-Z, a-z, 0-9) and the hyphen (-) character.
HA network interface	The *HA* interface (labeled *4* or *HA*) is an interface reserved for communication between the nodes of SSB clusters.
IETF-syslog protocol	The syslog-protocol standard developed by the Internet Engineering Task Force (IETF), described in RFC 5424 *The IETF syslog Protocol*.
key pair	A private key and its related public key. The private key is known only to the owner; the public key can be freely distributed. Information encrypted with the private key can only be decrypted using the public key.
LDAP	The Lightweight Directory Access Protocol (LDAP), is an application protocol for querying and modifying data using directory services running over TCP/IP.
log path	A combination of sources, filters, parsers, rewrite rules, and destinations: syslog-ng examines all messages arriving to the sources of the logpath and sends the messages matching all filters to the defined destinations.
logstore	A binary logfile format that can encrypt, sign, compress, and timestamp log messages.
log source host	A host or network device (including syslog-ng clients and relays) that sends logs to the syslog-ng Store Box. Log source hosts can be servers, routers, desktop computers, or other devices capable of sending syslog messages or running syslog-ng.
LSH	See *log source host*.

Management network interface	The *management* interface (labeled *2* or *MGMT*) is used exclusively for communication between SSB and the auditor or the administrator of the syslog-ng Store Box.
master node	The active SSB unit that is collecting the log messages when SSB is used in High Availability mode.
name server	A network computer storing the IP addresses corresponding to domain names.
node	An SSB unit running in High Availability mode.
output buffer	A part of the memory of the host where syslog-ng stores outgoing log messages if the destination cannot accept the messages immediately.
output queue	Messages from the output queue are sent to the target syslog-ng server. The syslog-ng application puts the outgoing messages directly into the output queue, unless the output queue is full. The output queue can hold 64 messages, this is a fixed value and cannot be modified.
overflow queue	See *output buffer*.
ping	A command that sends a message from a host to another host over a network to test connectivity and packet loss.
port	A number ranging from 1 to 65535 that identifies the destination application of the transmitted data. For example: SSH commonly uses port 22, web servers (HTTP) use port 80, and so on.
Public-key authentication	An authentication method that uses encryption key pairs to verify the identity of a user or a client.
redundant Heartbeat interface	A redundant Heartbeat interface is a virtual interface that uses an existing interface of the SSB device to detect that the other node of the SSB cluster is still available. The virtual interface is not used to synchronize data between the nodes, only Heartbeat messages are transferred.
regular expression	A regular expression is a string that describes or matches a set of strings. The syslog-ng application supports extended regular expressions (also called POSIX modern regular expressions).
relay mode	In relay mode, syslog-ng receives logs through the network from syslog-ng clients and forwards them to the central syslog-ng server using a network connection.
SSB	An abbreviation of the syslog-ng Store Box name.
slave node	The passive SSB unit that replaces the active unit (the master node) if the master becomes unavailable.
source	A way for SSB to receive syslog messages.

source, network	A source that receives log messages from a remote host using a network connection. The UDP, TCP, and TLS methods are supported using the BSD-syslog and the IETF-syslog protocols.
source, local	A source that receives log messages locally from SSB.
source driver	A communication method used to receive log messages.
SNMP	Simple Network Management Protocol (SNMP) is an industry standard protocol used for network management. SSB can receive SNMP messages from remote hosts and convert them to syslog messages, and can also send its own SNMP traps to a central SNMP server.
split brain	A split brain situation occurs when for some reason (for example the loss of connection between the nodes) both nodes of a SSB cluster become active (master). This might cause that new data (for example log messages) is created on both nodes without being replicated to the other node. Thus, it is likely in this situation that two diverging sets of data are created, which cannot be trivially merged.
syslog-ng	The syslog-ng application is a flexible and highly scalable system logging application, typically used to manage log messages and implement centralized logging.
syslog-ng agent	The syslog-ng agent for Windows is a log collector and forwarder application for the Microsoft Windows platform. It collects the log messages of the Windows-based host and forwards them to SSB using regular or SSL-encrypted TCP connections.
syslog-ng client	A host running syslog-ng in client mode.
syslog-ng Premium Edition	The syslog-ng Premium Edition is the commercial version of the open-source application. It offers additional features, like encrypted message transfer and an agent for Microsoft Windows platforms.
syslog-ng relay	A host running syslog-ng in relay mode.
syslog-ng server	A host running syslog-ng in server mode, like SSB.
TLS	Transport Layer Security (TLS) and its predecessor, Secure Sockets Layer (SSL), are cryptographic protocols which provide secure communications on the Internet. The syslog-ng application can encrypt the communication between the clients and the server using TLS to prevent unauthorized access to sensitive log messages.
template	A user-defined structure that can be used to restructure log messages or automatically generate file names.
traceroute	A command that shows all routing steps (the path of a message) between two hosts.

Index

Symbols

9 HA address, 234
<Initial window size>, 5
<name-of-the-logspace>, 61
<name>, 185
<Number of sources>, 5
[Set date and time separately], 135

A

AAA, 35, 68, 70-71, 73, 76, 78-82, 198, 200-202
Accept, 24
Access Control, 78-81, 103, 221
Access control, 146, 148, 159
accessing SSB using SSH, 104
Accounting, 82, 198, 200-201
Accounting settings, 82
accounting SSB, 82, 201
Activate sealed mode, 107
Activate slave, 89
Active Directory, 74
Active:, 33
Add, 19-20, 201
Add Chapter, 209
Add filter, 219
Add Subchapter, 209
Address, 38, 40-41, 76, 123, 158, 162, 164
Address/Netmask, 38
ADMIN, 110
Admin password, 26
admin password (see administrator password)
administrator password, 26
Administrator's e-mail, 25
Administrator's e-mail address, 43
Advanced, 19
Advanced mode, 136
After reboot, 96-99
Alert, 51-52
Alerting, 216
Alerting & Monitoring, 43, 45, 47-48, 50, 57, 62, 146, 148
alerts
  master, 51
  message rate, 50, 131
Alerts, 198, 203, 207, 211
alias interfaces, 36
alias IP addresses, 17, 20, 36
All, 140, 197
All messages in one file, 145, 147
All Tasks, 116
Allow compression, 130
Allow private key to be exported, 119
Allowed, 120
Allowed group, 151, 154
Always, 51-52
Anonymous, 65
Append, 124
Application Policies, 119
Apply, 121, 193-194
Archive & Cleanup, 198, 204
Archive now, 62
Archive target, 204
Archive/Cleanup, 150
Archive/Cleanup policies, 59
Archive/Cleanup policy, 61-62, 146, 148
archiving, 59, 100
  log messages, 62
  NetApp devices, 56, 60, 64
  notifications, 204
  Windows 2008 R2, 57, 61, 65
artificial ignorance, 211
  message classification, 217
Artificial ignorance, 212
Asynchronous data replication, 89
auditing configuration changes, 82
auth-view, 81
auth-write, 81
Auth. method, 46
Auth. password, 46
Authenticate as client, 75
Authenticated Users, 120
authentication, 7-8
  LDAP, 68
  to Microsoft Active Directory, 68, 72
  to RADIUS servers, 76
  users, 68, 72
Authentication key, 63
Authentication method, 45, 70, 76
Authentication password, 45
Authentication settings, 73
Author, 83, 200, 202
Authorized keys, 106

Autoclose successful commit messages, 34
Automatically start archiving, 50
Available columns, 201
Available dynamic columns, 184-185
Available static columns, 184

B

Back, 22
Back to Main menu, 107
Backup, 58, 149
Backup & Archive/Cleanup, 55, 59, 236
Backup ALL, 58
Backup All, 149
Backup now, 58
Backup policies, 55
Backup policy, 58, 145, 148
backups, **55**
  encrypting, 58
  file ownerships, 67
  file permissions, 67
  NetApp devices, 56, 60, 64
  notifications, 204
  restore, 235
  Windows 2008 R2, 57, 61, 65
Bar chart, 195
Base DN, 74
Base-64 encoded X.509 (.CER), 117
Basic, 40, 59
Basic mode, 133
Basic Settings, 31, 33-35, 37, 39, 41-45, 47-50, 57-58, 81, 84-85, 89-90, 92, 96-102, 105-107, 110, 112, 114, 122-123, 125, 146, 148, 153, 177, 204-206, 209, 222-224, 226-229, 232-234, 236, 240, 246
Basic settings, 38, 198-199
basic-view, 81
basic-write, 81
Bind DN, 74
Bind Password, 74
Blink system identification lights, 222
Boot, 99
Boot firmware, 98
Boot firmware version, 87
Boot firmwares, 96-97
Boot options, 245
Boot shell, 109, 232
Broken, 228
Browse, 22, 24, 59, 75, 100, 102, 114, 124, 144, 177-178, 216, 236
browser requirements, 31
browsers, 31
  supported versions, 31
browsing log messages, 181
browsing reports, 205
built-in sources, 126
Button, xiii

C

CA certificate, 111
CA X.509 certificate, 75, 122
center, 166
Certificate, 114, 144
Certificate >, 191, 193
Certificate Authorities, 178
Certificate Template to Issue, 121
Certificate Templates, 118
certificates, 7
  accepted formats, 113
  changing, 110
  extendedKeyUsage, 114-115
  for TLS authentication, 177
  LDAP servers, 75
  managing, 110
  Timestamping Authority, 110
  TSA, 115
  uploading, 113
  web interface, 27, 110
  X509v3 Extended Key Usage, 114-115
Change, 192
Change root password, 106
changelog, 82
changelogs, 82, 201
changing certificates
  Timestamping Authority, 110
CIFS, 151, 154
Class, 216
classifier_class, 219
classifier_rule_id, 219
classifying messages, 211
  alerts, 203
  pattern matching concepts, 213
cleanup, 59
Cleanup if unchanged for, 205
Cleanup now, 205
Clear all filters, 200
Clear all statistics, 204
Client address, 46

Client key, 75
Client X.509 certificate, 75
Collect and save current system state info, 224
collecting debug information, 224
collecting system-state information, 224
commit log, 82
Community, 44, 46, 127, 164
Completing the Certificate Export Wizard, 118
compliance, 2
Compressed logstore, 145
configuration
  backups, 55
  changes, 82
  delimiters, 145
  log paths, 166
  logspaces, 143
  network interfaces, 36
  sources, 128, 131
Configuration, 207-208
Configuration changed, 52
Configuration changes, 207
Configuration saved successfully, 34
Configure, 110
Confirm password, 101, 106
Connected, 229
Connected (Disk Failure), 229
Connected syslog peers, 226
consistent, 228
console menu, 103
Content, 31
Content Advisor, 31
controlling SSB
  rebooting, 84
  shutting down, 84
Convert to Cluster, 228-229, 240
Converted, 229
converting SNMP to syslog, 127
Copy to File, 117
Copy-paste, 75
Copy-paste key, 177-178
Core files, 223
core files, 223
Core firmware, 98-99
Core firmwares, 96-97
Core shell, 21, 239
Core Shell, 232
Country, 27, 112
CPU, 227
Cracklib check on password, 77
Create new ruleset, 214
creating local destinations, 143
creating log spaces on volumes, 140, 143
creating logspaces, 143
creating sources, 128, 131
CRL URL, 179
CSV Export, 184
Current, 98
Current master, 86
Custom, 145, 148, 161
Custom address, 210
Custom columns, 159
custom log files, 143
Custom message part only, 163
Custom on-wire message, 163
custom reports, 208
custom sources
  SQL, 131
  syslog, 128
Customer, 99
Customize columns, 184, 195
Customize Columns, 200

D

Daily reports, 206
Dashboard, 31, 34, 199, 204-205, 226-227
dashboard, 204, 226
Dashboard Statistics, 204, 226
data
  backups, 55
Data and configuration backup failed, 53
Data archiving failed, 53
Database error occurred, 53
database format, 160
Database name, 159
Database Server, 157
database templates, 160
Database type, 133, 158
Date & Time, 40
Date & Time Settings, 40
date and time, 40
  configuring, 41
Day:, 224
debug logging, 224
Debug logging, 224
Decryption private keys >, 194
Default gateway, 26

default reports, 207
default sources, 126
default usergroups, 81
Degraded, 228
DEGRADED, 230
Degraded (Disk Failure), 228
Degraded Sync, 229, 233
DEGRADED-WORKING, 230
deleting
  log files, 60, 149
Delimiters, 140
Description, 202
Destination, 168
destinations, 3, 140, 157
Destinations, 50, 157, 159, 161, 164, 168
Details, 117
Directory name, 204
Disable, 31, 107, 232
Disabled, DES or AES, 45-46
disabling message parsing, 131
Disk, 227
Disk space fill up prevention, 49, 54
Disk usage is above the defined ratio, 50, 54
Disk utilization maximum, 49
Displayed columns, 184-185
Distinguished name, 179
DNS
  server, 39
DNS Cache expiry, 176
DNS search domain, 39
DNS server, 25
Do not parse messages, 131
Domain, 66, 152
Domain controller, 153
Domain mode, 152
Domain name, 25
Don't parse messages, xv
downgrading the firmware
  rollback, 98
Download, 206, 224
Download MIBs, 47
DRBD
  adjusting synchronization speed, 89
DRBD asynchronous mode, 88-89
DRBD status, 86-87, 89, 228-229
  connected, 229
  Connected (disk failure), 229
  Invalidated, 229
  split brain, 230-231
  Sync source, 229
  Sync target, 229
  wfconnection, 230
DRBD sync rate limit, 86-87, 89, 240
Duplicate Template, 118

E

e-mail alerts, 41-42, 47-48
e-mailing reports, 42
Empty, 150
Enable, 107, 145, 169
Enable cracklib, 71
Enable debug logs, 224
Enable management interface, 38
Enable nested groups, 73
Enable password authentication, 105
Enable remote SSH access, 105
Enable statistics for, 205
Enabled, 204
Encoding, 131
Encrypt configuration, 59
Encrypt with password, 101
encrypting log messages, 7
Encryption, 74
Encryption certificate, 144
Encryption method, 45-46
Encryption passphrase, 23
Encryption password, 45-46, 101-102
Engine ID, 44
Enhanced Key Usage, 115, 117
Enroll, 120
Ethernet links, 38, 246
Export, 66-67, 101, 216
Export all to CSV, 195
Export as CSV, 196, 200
Export configuration, 57, 80, 101
Export ruleset, 216
exporting
  search results, 200
  SSB configuration, 100
exporting pattern database, 216
External interface, 36, 227
External interface — IP address, 25
External interface — Netmask, 26
external timestamps, 175

F

facilities, 12, 14
Facility, 135, 137, 140
Failed DNS cache expiry, 176
Fast follow mode, 135, 137
feature releases, 9
Fetch data in every X seconds, 135, 137
Field name, 83, 202
file destinations, 140
Filename template, 145, 147
Filter, 168
Filter ACLs, 80
filtering messages, 169
filtering search results, 200
filters, 4, 169
Final, 166, 168
finding patterns, 214
Fingerprint, 203
Finish, 29-30, 118
Finishing the Setup, 243
firmware, 9
  high availability, 9
  rollback, 98
  update, 95, 97
Firmware, 80
Firmware is tainted, 53
Flow, 5, 168
flow-control, 5
  multiple destinations, 7
Flush lines, 159
Force BIOS Setup, 245
Freeze, 201
Full, 160
Full domain name, 153

G

Gateway, 38, 40
General alert, 53
General error, 53
Generate, 27-28, 63
Generate All, 113
Generate new self-signed certificate, 27
Generate partial daily report, 208
Generate partial monthly report, 208
Generate partial weekly report, 208
Generate reports for today, 207
Generate Server certificate, 113
Generate this report every, 209
Generate time, 207
Generate TSA certificate, 113
Generated reports, 205, 208
Generation 1, 246
Get current size, 149
Global alerts, 51
Global master alert, 52
GPG, 58
GPG encryption, 101
Grant access for the following user groups, 197
Group, 80, 159
Group Management, 71, 81
group management
  local, 71
GroupOfUniqueNames membership attribute name, 75
Groups, 69

H

HA, 26, 90, 228 (see High Availability)
HA address, 26, 38
HA interface, 38
HA link speed, 87
HA MAC, 90
HA node state changed, 53
HA UUID, 86
HA:, 34
Half, 228
Hardware error occured, 53
Hardware informantion, 222
Hardware serial number, 108
hardware serial number, 108
Heartbeat, 233
High Availability, 9, 85, 89-90, 92, 107, 228-229, 233-234, 240
  address, 26, 38
  adjusting synchronization speed, 89
  installation, 240
  log messages, 87
  manual takeover, 89
  next-hop monitoring, 92
  Node HA status, 228
  Node HA UUID, 228
  node replacement, 233
  reboot cluster, 88
  recovery, 230
  redundant heartbeat interfaces, 90

status, 228
synchronizing time, 41
High availability, 97-98
high availability mode, 85
history of changes, 201
Host, 135, 140, 145
Host limit, 100
Host:, 32
Hostlist, 127, 131, 152, 154
hostlists, 122, 127, 152, 154
changing, 123, 125
creating new, 122
importing, 123
modifying, 123, 125
Hostlists, 122, 124
Hostname, 25, 39, 153, 202-203, 223
Hosts:, 33

I

Idle time before destination is closed, 174
Ignore, 123
Ignore ambiguous program field, 131
ILOM, 107
Import, 23, 236
Import configuration, 80, 102, 236
Import from file, 124
importing
certificates, 113
SSB configuration, 102
importing certificates, 177
importing pattern database, 216
Include file list, 57
Indexed, 161
Indexed fields, 140
Indexer, 145
indexing
delimiters, 145
Install a new SCB, 242
Installation Steps, 242
Installer, 242
installing
SSB, 238
Integrated Lights Out Management, 107
Intelligent Platform Management Interface, 107
Interface IP, 91-92
Interfaces, 37-39
Interfaces for Heartbeat, 87, 90
Internet Options, 31
Internet Protocol (TCP/IP), 18
Interval, 207
INVALID, 230
Invalidated, 229
IP Address, 20
IP Addresses, 19
IP Settings, 19
IPMI, 107
IPMI default gateway IP, 108
IPMI IP address, 108
IPMI IP address source, 108
IPMI subnet mask, 108
ISO date, 163

J

JavaScript, 31
Join domain, 153
Join HA, 228-229, 233
Jump to, 199
Jump to last option, 183

K

Key, 59, 114, 177-178
Key >, 192, 194
Key Usage, 119

L

Label, xiii
Last login:, 32
LDAP, 73
LDAP authentication, 68, **72**
LDAP groups
nested groups, 73
LDAP servers
certificates, 75
custom attributes, 75
failover, 73
GroupOfUniqueNames membership attribute name, 75
Microsoft Active Directory on Windows 2000, 75
mutual authentication, 75
POSIX group membership attribute name, 75
Username (userid) attribute name, 75
Windows 2000, 75
Least, 195, 197
Legacy, 130, 160, 163
license, 10

update, 99
License, 99-100
License limit reached, 53
License:, 33
Limit of alerts sent out in a batch, 52
Listening address, 129
Load 15:, 34
Load 1:, 34
Load 1|5|15 maximum, 49
Load average, 227
Local, 70, 175
local, 166, 198
Local Area Connection, 17
local console, 103
local name resolution, 176
local SSB users, 68
local time, 12, 14
Local Users, 68, 81
local users
password management, 69
usergroups, 71
Locality, 27, 113
lock management, 35
Locked:, 33
Log, 6, 50-52, 58, 61-62, 82, 103, 127-129, 132, 143, 146, 149, 152, 154, 157, 159, 161, 164, 166-168, 170, 173, 175-177, 194, 204, 211-212, 214, 216, 219, 226, 236
log
debug mode, 224
system state, 224
tailing, 223
viewing, 223
Log Alerts, 223
log message structure
BSD-syslog protocol, 10
IETF-syslog protocol, 12
legacy-syslog protocol, 10
RFC 3164, 10
RFC 5424, 12
log messages
alerts, 203
browsing, 181
reports, 205
searching, 181
structure of, 10
log paths, 3
creating new, 166
defaults, 166
flow-control, 5
log space size, 149
log spaces
accessing shares, 155
deleting log files, 149
instant archiving, 149
instant backup, 149
logstore limitations, 141
manual backup, 149
sharing, 150
size of, 149
using logcat, 142
log statements (see log paths)
log-view, 82
log-write, 82
logcat, 142
logging procedure, 4
Login failed, 52
Logout, 22, 107, 239
Logout from the management interface, 52
Logs, 35, 103, 140
logspace, 140
Logspace, 204
Logspace exceeded warning size, 146, 148
Logspace name, 181
Logspaces, 227
logstore, 140, 142
browsing encrypted, 190
cipher method used for encrypting, 175
digest method used for encrypting, 175
limitations of, 141
timestamping authority, 175
using logcat, 142
LogStore, 144
Logtype, 224
Long Term Supported releases, 9
LTS releases, 9

M

Mail settings, 42, 49, 205-206, 209
Main menu, 31
Make HA IP permanent, 229
Make this extension critical, 119
Manage, 118
Management, 33, 35, 38, 41-42, 44-45, 47, 49, 57-59, 61, 101, 105-106, 110, 112, 114, 122, 146, 148, 177, 205-206, 209, 223-224, 236

Management enabled, 36, 38
Management Information Base, 47
Management interface, 37-38, 227
managing certificates
  Timestamping Authority, 110
Manual archiving, 204
Master alert, 51
Match, 123
Maximum, 47, 50, 52
Maximum connections, 129
Maximum number of files in notification, 57
Maximum number of search results, 145
Maximum number of statistics to process, 204
MD5 or SHA1, 45-46
Memory, 227
Memory buffer size, 147
Memory limit, 140
Menu, xiii
Message, 83, 140, 145, 204, 213, 216
message classification, 211
  alerts, 203
message destinations, 140, 157
Message dialog, 35
message facilities, 12, 14
message filtering
  using parsers, 218
message filters, 169
Message part, 171
message rate alerting, 50, 131
Message rate alerting, 50, 135, 137
Message rate alerting statistics, 51-52
Message size, 11
message sources, 126
Message throttle, 163
Message:, 224
Messages fetched in a single poll, 5-6
MIB, 47
Microsoft Active Directory
  supported platforms, 75
Microsoft Windows Network, 62
Minimal password strength, 71, 77
Minimum, 50, 52
Modify, 110
Modify User, 110
Modules:, 33
monitoring, 42-43, **47**, 48
  MIB, 47
Month, 226
Monthly reports, 206
mounting shares, 155
mutual authentication, 8

N

Name, 206, 215
name resolution, 176
  local, 176
Name resolving, 176
Name/value pairs, 140, 145
Naming, 39
NetApp, 56, 60, 64
Netmask, 20, 38, 40
Network, 36-37, 39, 246
Network Connections, 17
Network connections, 227
network interfaces
  alias interfaces, 36
  alias IP addresses, 36
  configuring, 36
  configuring interface speed, 37
  external interface, 8, 16, 36
  HA interface, 8
  IPMI interface, 9
  management interface, 8, 37, 39
network shares, 150
Networks, 38
New, 121
New root password, 106
New value, 83, 202
New Virtual Machine Wizard, 247
Next, 22, 24, 26-27, 117
Next hop IP, 93
Next hop monitoring, 87, 92, 233
next-hop router monitoring, 92
NFS, 57, 61, 66, 150-151, 154-155
Nick name, 39
nickname, 39
No, 243
No archiving, 60
No encryption, 101
no-parse, 131
Node HA state, 87
Node HA status, 228
Node HA UUID, 87, 228
Node ID, 87
NOT USED, 230
NTP server, 25

NTP servers, 41
number of active hosts, 33
number of active senders, 33
Number of entries, 197
Number of passwords to remember, 77

O

OK, 119, 121, 201, 209, 230
Old value, 83, 202
On, 222
Once, 51-52
Only accept certificates authenticated by the specified CA certificate, 75
Only cleanup, no archiving, 60
Only from persistent configuration, 176
Only with the name, 140
operational-report, 207
Options, 6, 51-52, 127, 129, 173, 175-177, 204, 211-212, 216, 226, 245
Organization, 27, 113
Organization unit, 28, 113
Other node, 84-85, 91, 93, 97, 233-234
out-of-band management, 107
output buffer, 6
Output disk buffer, 159, 164-165
Output memory buffer, 5, 159, 164-165

P

Page, 83, 202
parameters
  keep_hostname() , 128
  keep_timestamp() , 128
  log_fetch_limit() , 5
  log_fifo_size() , 5
  log_iw_size() , 6
  max_connections() , 128
parsing messages, 217
  filtering parsed messages, 218
partial reports, 207
password
  admin, 26
  changing the root, 106
  root, 26, 105
Password, 65, 69, 153, 158
Password expiration, 70, 77
password policies, 69
Password provided by database, 70
Path, 64
Paths, 103, 166-167, 170, 219
Pattern, 216
pattern database, 211
  adding patterns, 214
  browsing, 214
  create ruleset, 214
  creating parsers, 217
  export database, 216
  export ruleset, 216
  import database, 216
  import ruleset, 216
  structure of, 212
  using the results, 218
Pattern Database, 214, 216
pattern databases
  pattern matching precedence, 213
pattern matching, 211
  procedure of, 213
  searching for patterns, 214
patterndb (see pattern database)
Peer configuration, 207
Peer configuration change, 198
Peer Configuration Change, 202
Peer verification, 129
Pending Requests, 116
Per application, 145, 148
Per host, 145, 148
Per host and application, 145, 148
Permanent >, 191
Persistent hostname list, 176
persistent name resolution, 176
Pid, 140
Pie chart and List, 195
Ping, 153
ping, 222
Ping host, 153
Policies, 55, 59, 82, 122, 124, 150, 152, 236
policies-view, 82
policies-write, 82
Policy, 204
Port, 64, 158, 162, 164, 223
Posix, 74
POSIX group membership attribute name, 75
Preferences, 31, 34
preferences, 32
preventing message loss (see flow-control)
Primary DNS server, 39
Priority, 140

private key
accepted formats, 29, 113
Private keystore, 191, 193
Processes, 227
Production MAC, 90
Program, 135, 140, 145, 203, 213, 216
Program pattern, 215-216
Properties, 17-18
public-key authentication on SSB, 104
Put all columns into SDATA, 135, 137

Q

Query, 63
querying SSB via SNMP, 45

R

RADIUS, 76
RADIUS authentication, 76
Raid status, 87
Raid status changed, 53
Raid status:, 33
Rate limit, 127-128
Read old records, 135, 137
reboot, 84
Reboot, 96-97, 99, 234, 243
Reboot Cluster, 88, 92-93, 240
Reboot cluster, 89
Recipient, 210
Recommended, 112
Redundant, 90, 228
redundant heartbeat interfaces, 90
Redundant Heartbeat status, 90, 228
redundant heartbeat status
degraded, 230
degraded-working, 230
invalid, 230
not used, 230
ok, 230
releases, 9
Remaining time:, 33
Remote, 175
remote destinations, 157
remote servers, 161
SNMP, 164
SQL databases, 157
Remote host, 161
remote server, 157
Remove, 201
Replace, 124
Replacement value, 171
report, 82
Report from, 207
Report settings, 197
Report subchapter name, 197
Report to, 207
Reporting, 35, 199, 209
reports, 205, 207
browsing, 205
contents, 207
custom reports, 208
default, 207
e-mailing, 42
partial reports, 207
Reports, 82, 199, 205, 207-209
Reports are accessible by the following groups, 209
Require commit log, 82-83, 202
REST, 221
restart, 84
Restart syslog-ng, 100, 123, 125
Restore, 149
Restore ALL, 236
Restore now, 236
restoring a backup, **235**
Retention time, 159
Retention time in days, 60-61
Revert Configuration, 98
reverting the firmware version, 98
Rewrites, 170
RLTP, 129-130
RLTP TLS, 129-130
Root password, 26
root password, 26, 105-106
Routing table, 38-39
management traffic, 39
RPC API, 221
client requirements, 221
documentation, 221
requirements, 221
Rsync over SSH, 56, 62
rsync over SSH, 62
Rule description, 204
Rule ID, 204
Rules, 216
Ruleset name, 213, 216
Run, 116

S

samba shares, 150, 155
Sampling interval, 204, 226
Save, 78-79
Save As Report subchapter, 197
Save the collected debug info, 225
Scale, 199
Seal the box, 27
sealed mode, 107
Sealed mode, 107
Search, 35, 79-80, 82, 103, 181, 184, 198, 202-204, 213, 216, 221, 223, 227
search, 82
  boolean, 187
  wildcard, 187
Search in, 227
search results, 182
  statistics, 195
searching log messages, 181
Secondary DNS server, 39
Security passphrase, 192
Select resolution, 227
Send e-mail alerts to, 43
Send e-mails as, 43
Send even empty reports, 210
Send notification on all events, 57, 61
Send notification on errors only, 57, 61
Send reports in e-mail, 210
Send reports to, 43, 205-206, 209
Sender address, 203
Senders:, 33
Serial, 100
serial number of SSB, 108
Server Address, 73-75
Server Authentication, 119
Server certificate, 110-111
Server host key, 63
Server private key, 27
Server URL, 175
Server X.509 certificate, 27
Service control, 100, 123, 125, 232
Set, 59, 75, 114
Set Date & Time, 40
Set Default Port, 133
Settings, 17, 70, 73, 76, 82, 202, 247
Severity, 135, 137
SHA-1 fingerprint, 179
Share, 66
Shared secret, 77
Shares, 150, 152
sharing log files, 150, 155
Sharing policy, 146, 148, 152, 154
Shells, 21, 109, 232, 239
Show, 213, 216
shutdown, 84
Shutdown, 97
Shutdown, Suspend, Reset, 245
Signature, 203
Signature is proof of origin, 119
Simple Network Management Protocol, 43, 47
Size, 149
size of a log space, 149
SMB/CIFS, 56, 60, 62, 64-65
SMB/CIFS options, 150, 152
SMTP server, 25, 42
SMTP server address, 42
SNMP
  alerts, 43, 47-48
  messages, 127
  queries, 45
  server, 43
  SSB MIB, 47
SNMP agent settings, 45
SNMP destination, 164
SNMP server address, 44
SNMP settings, 49
SNMP source, 127
SNMP trap settings, 44
SNMP v2c, 44, 164
SNMP v2c agent, 46
SNMP v3, 44, 164
SNMP v3 agent, 46
SNMPv3, 45
Source, 167
sources, 3, 126
  creating new, 128, 131
  defaults, 126
  SNMP, 127
Sources, 50, 128, 132, 176
Spaces, 50, 58, 61-62, 103, 140, 143, 146, 149, 152, 154, 168, 181, 194, 198, 236
spaces, 140
  creating new, 143
  indexer delimiters, 145
Speed, 37
Split brain, 229-230

split brain, 229-231
Spoof source address, 162
SQL, 132
sql sources
  customized queries, 136
  variables, 136
SQL templates, 160
SSB
  accounting, 82
  administrators, 68
  certificate, 110
  changelogs, 201
  configuration (see SSB configuration)
  configuration changes, 82
  exporting the configuration of, 100
  hostname, 39
  importing the configuration of, 102
  installation, 238
  logs, 82
  nickname, 39
  reports, 205
  web certificate, 27
SSB configuration
  exporting, 100
  importing, 102
SSB options, 173
SSH
  console, 103
SSH connections
  accessing SSB, 104
SSH server on SSB, 104
SSH settings, 105
SSL certificate, 110, 112, 114, 122
SSL certificates, 177
SSL/TLS, 74
stable releases, 9
Standalone, 228
Standalone mode, 150
Start, 116, 225
Start menu, 17
Start time, 56, 60
STARTTLS, 74
State or Province, 28, 113
statistics, 195
  settings, 204
  time-based, 204
  top-least, 204
Statistics, 197
Status, 86, 228, 233-234
status history, 226
status information via SNMP, 45
Stop, 225
Submit new request..., 116
Successful login, 52
supported browsers, 1, 31
supported timestamping protocols, 175
Suppress timeout, 163
Swap, 202
Swap utilization maximum, 49
Sync Master, 41
Sync now, 41
Sync Slave to Master, 41
Sync source, 229
Sync target, 229
synchronizing data
  adjusting synchronization speed, 89
SyncSource, 233
SyncTarget, 233
Syslog, 129-130
Syslog flags, 131
Syslog protocol, 130, 163
Syslog traffic, indexing & search:, 100, 123, 125, 232
syslog-ng, 226
  certifications, 177
  logging configuration changes, 202
  options, 173
syslog-ng options, 173
Syslog-ng statistics, 204-205
syslog-ng statistics, 227
syslog-ng traffic statistics, 207
System, 57, 59, 84, 96-102, 107, 123, 125, 222, 232-233, 236
System backup, 58-59, 101, 236
System backup policy, 57-58
System contact, 46
System control, 84, 99
System Control, 96
System debug, 224
System description, 46
System health information, 207
System location, 46
System monitor, 31, 33, 229
system monitor
  number of active hosts, 33
  number of active senders, 33
System Monitor, 35

System related traps, 50
system statistics, 226

T

Table, 134
Table of contents, 209
Table rotation, 159
Tags, 140, 195
Tail, 224
Target server, 56, 61
Target settings, 61-62, 65
TCP, 129
Template, 145, 147-148
Template display name, 118
Temporary >, 193
Test, 43, 76
Test connection, 160
Test connection and fetch tables, 133
Test data retrieving, 136-137
Text file, 147
This node, 84-85, 91, 93, 96-97, 99
Time Stamping, 119
Time sync lost, 53
time synchronization, 41
  in HA mode, 41
Time-based statistics, 204-205
Time:, 33
timeout
  web session, 35
timestamp, 12, 14
Timestamp, 83, 202-204
Timestamp fractions of a second, 159, 164-165
Timestamping Authority
  certificate of, 110
Timestamping error occured, 53
Timestamping frequency, 145
timestamping OID, 175
timestamping protocol, 175
timestamping server, 175
Timezone, 25, 41, 131, 135, 137, 159, 164-165
TLS, 7, 129
TLS certificate, 177
TLS private key, 177-178
TLS settings, 129, 177
Tools, 31
Top, 197
Top/Least statistics, 204
traceroute, 222
tracking configuration changes, 82
Transport, 129, 162
transport layer security (see TLS)
Troubleshooting, 34, 81, 153, 198, 222-224
troubleshooting, 222
Trusted, 129
Trusted distinguished names, 179
Trusted fingerprints, 179
TSA certificate, 110-111
TSA private key, 122
TSA X.509 certificate, 122
Type, 74, 79-80, 144, 147, 151, 154

U

UDP, 129
Unblock Slave Node, 88
Unique ID column, 134
update
  firmware, 95
  in high availability, 97
  license, 99
upgrade
  license, 99
Upload, 24, 59, 75, 100, 102, 114, 122, 124, 144, 177-178, 216
Upload key, 177-178
uploading certificates, 113
Use DNS, 128, 131, 176
Use FQDN, 131
use static subchapters, 208
User, 69
User database, 70
user groups, 68
User info, 31
user management, 81
  creating usergroups, 79
  finding privileges, 80
  modifying usergroup privileges, 78
  naming usergroups, 80
  searching usergroups, 80
User menu, 31-32, 34
User Menu, 191, 193
user preferences, 32
User:, 32
usergroups
  local, 71
Username, 44, 46, 63, 65, 153, 158
Username (userid) attribute name, 75

users
web interface, 68, 72
Users, 110

V

Validity, 100, 203
variables
in sql queries, 136
Verify password, 69
Version, 203
Version details, 97, 233
View, 224
View graph, 227
View log files, 223
Visible columns, 185, 201
Visualization, 197
VMware Tools, 245
volumes, 140, 143

W

Warning size, 146, 148
Warning, all data on the hard drive(s) will be erased.
Are you sure?, 242
web browsers, 31
Web interface timeout, 33, 35
Web Server, 118
web session timeout, 35
Week, 226
Weekly reports, 206
Welcome Wizard, 22
WFConnection, 230
Windows Certificate Authority, 115

Y

Year, 226
Yes, 242-243

List of SSB web interface labels

Symbols

9 HA address, 234
<Initial window size>, 5
<name-of-the-logspace>, 61
<name>, 185
<Number of sources>, 5
[Set date and time separately], 135

A

AAA, 35, 68, 70-71, 73, 76, 78-82, 198, 200-202
Accept, 24
Access Control, 78-81, 103, 221
Access control, 146, 148, 159
Accounting, 82, 198, 200-201
Accounting settings, 82
Activate sealed mode, 107
Activate slave, 89
Active Directory, 74
Active:, 33
Add, 19-20, 201
Add Chapter, 209
Add filter, 219
Add Subchapter, 209
Address, 38, 40-41, 76, 123, 158, 162, 164
Address/Netmask, 38
ADMIN, 110
Admin password, 26
Administrator's e-mail, 25
Administrator's e-mail address, 43
Advanced, 19
Advanced mode, 136
After reboot, 96-99
Alert, 51-52
Alerting, 216
Alerting & Monitoring, 43, 45, 47-48, 50, 57, 62, 146, 148
Alerts, 198, 203, 207, 211
All, 140, 197
All messages in one file, 145, 147
All Tasks, 116
Allow compression, 130
Allow private key to be exported, 119
Allowed, 120
Allowed group, 151, 154
Always, 51-52
Anonymous, 65
Append, 124
Application Policies, 119
Apply, 121, 193-194
Archive & Cleanup, 198, 204
Archive now, 62
Archive target, 204
Archive/Cleanup, 150
Archive/Cleanup policies, 59
Archive/Cleanup policy, 61-62, 146, 148
Artificial ignorance, 212
auth-view, 81
auth-write, 81
Auth. method, 46
Auth. password, 46
Authenticate as client, 75
Authenticated Users, 120
Authentication key, 63
Authentication method, 45, 70, 76
Authentication password, 45
Authentication settings, 73
Author, 83, 200, 202
Authorized keys, 106
Autoclose successful commit messages, 34
Automatically start archiving, 50
Available columns, 201
Available dynamic columns, 184-185
Available static columns, 184

B

Back, 22
Back to Main menu, 107
Backup, 58, 149
Backup & Archive/Cleanup, 55, 59, 236
Backup ALL, 58
Backup All, 149
Backup now, 58
Backup policies, 55
Backup policy, 58, 145, 148
Bar chart, 195
Base DN, 74
Base-64 encoded X.509 (.CER), 117
Basic, 40, 59
Basic mode, 133

Basic Settings, 31, 33-35, 37, 39, 41-45, 47-50, 57-58, 81, 84-85, 89-90, 92, 96-102, 105-107, 110, 112, 114, 122-123, 125, 146, 148, 153, 177, 204-206, 209, 222-224, 226-229, 232-234, 236, 240, 246
Basic settings, 38, 198-199
basic-view, 81
basic-write, 81
Bind DN, 74
Bind Password, 74
Blink system identification lights, 222
Boot, 99
Boot firmware, 98
Boot firmware version, 87
Boot firmwares, 96-97
Boot options, 245
Boot shell, 109, 232
Broken, 228
Browse, 22, 24, 59, 75, 100, 102, 114, 124, 144, 177-178, 216, 236
Button, xiii

C

CA certificate, 111
CA X.509 certificate, 75, 122
center, 166
Certificate, 114, 144
Certificate >, 191, 193
Certificate Authorities, 178
Certificate Template to Issue, 121
Certificate Templates, 118
Change, 192
Change root password, 106
changelog, 82
CIFS, 151, 154
Class, 216
classifier_class, 219
classifier_rule_id, 219
Cleanup if unchanged for, 205
Cleanup now, 205
Clear all filters, 200
Clear all statistics, 204
Client address, 46
Client key, 75
Client X.509 certificate, 75
Collect and save current system state info, 224
Community, 44, 46, 127, 164
Completing the Certificate Export Wizard, 118
Compressed logstore, 145
Configuration, 207-208
Configuration changed, 52
Configuration changes, 207
Configuration saved successfully, 34
Configure, 110
Confirm password, 101, 106
Connected, 229
Connected (Disk Failure), 229
Connected syslog peers, 226
consistent, 228
Content, 31
Content Advisor, 31
Convert to Cluster, 228-229, 240
Converted, 229
Copy to File, 117
Copy-paste, 75
Copy-paste key, 177-178
Core files, 223
Core firmware, 98-99
Core firmwares, 96-97
Core shell, 21, 239
Core Shell, 232
Country, 27, 112
CPU, 227
Cracklib check on password, 77
Create new ruleset, 214
CRL URL, 179
CSV Export, 184
Current, 98
Current master, 86
Custom, 145, 148, 161
Custom address, 210
Custom columns, 159
Custom message part only, 163
Custom on-wire message, 163
Customer, 99
Customize columns, 184, 195
Customize Columns, 200

D

Daily reports, 206
Dashboard, 31, 34, 199, 204-205, 226-227
Dashboard Statistics, 204, 226
Data and configuration backup failed, 53
Data archiving failed, 53
Database error occurred, 53
Database name, 159
Database Server, 157

Database type, 133, 158
Date & Time, 40
Date & Time Settings, 40
Day:, 224
Debug logging, 224
Decryption private keys >, 194
Default gateway, 26
Degraded, 228
DEGRADED, 230
Degraded (Disk Failure), 228
Degraded Sync, 229, 233
DEGRADED-WORKING, 230
Delimiters, 140
Description, 202
Destination, 168
Destinations, 50, 157, 159, 161, 164, 168
Details, 117
Directory name, 204
Disable, 31, 107, 232
Disabled, DES or AES, 45-46
Disk, 227
Disk space fill up prevention, 49, 54
Disk usage is above the defined ratio, 50, 54
Disk utilization maximum, 49
Displayed columns, 184-185
Distinguished name, 179
DNS Cache expiry, 176
DNS search domain, 39
DNS server, 25
Do not parse messages, 131
Domain, 66, 152
Domain controller, 153
Domain mode, 152
Domain name, 25
Don't parse messages, xv
Download, 206, 224
Download MIBs, 47
DRBD asynchronous mode, 88-89
DRBD status, 86-87, 89, 228-229
DRBD sync rate limit, 86-87, 89, 240
Duplicate Template, 118

E

Empty, 150
Enable, 107, 145, 169
Enable cracklib, 71
Enable debug logs, 224
Enable management interface, 38
Enable nested groups, 73
Enable password authentication, 105
Enable remote SSH access, 105
Enable statistics for, 205
Enabled, 204
Encoding, 131
Encrypt configuration, 59
Encrypt with password, 101
Encryption, 74
Encryption certificate, 144
Encryption method, 45-46
Encryption passphrase, 23
Encryption password, 45-46, 101-102
Engine ID, 44
Enhanced Key Usage, 115, 117
Enroll, 120
Ethernet links, 38, 246
Export, 66-67, 101, 216
Export all to CSV, 195
Export as CSV, 196, 200
Export configuration, 57, 80, 101
Export ruleset, 216
External interface, 36, 227
External interface — IP address, 25
External interface — Netmask, 26

F

Facility, 135, 137, 140
Failed DNS cache expiry, 176
Fast follow mode, 135, 137
Fetch data in every X seconds, 135, 137
Field name, 83, 202
Filename template, 145, 147
Filter, 168
Filter ACLs, 80
Final, 166, 168
Fingerprint, 203
Finish, 29-30, 118
Finishing the Setup, 243
Firmware, 80
Firmware is tainted, 53
Flow, 5, 168
Flush lines, 159
Force BIOS Setup, 245
Freeze, 201
Full, 160
Full domain name, 153

G

Gateway, 38, 40
General alert, 53
General error, 53
Generate, 27-28, 63
Generate All, 113
Generate new self-signed certificate, 27
Generate partial daily report, 208
Generate partial monthly report, 208
Generate partial weekly report, 208
Generate reports for today, 207
Generate Server certificate, 113
Generate this report every, 209
Generate time, 207
Generate TSA certificate, 113
Generated reports, 205, 208
Generation 1, 246
Get current size, 149
Global alerts, 51
Global master alert, 52
GPG encryption, 101
Grant access for the following user groups, 197
Group, 80, 159
Group Management, 71, 81
GroupOfUniqueNames membership attribute name, 75
Groups, 69

H

HA, 26, 90, 228
HA address, 26, 38
HA interface, 38
HA link speed, 87
HA MAC, 90
HA node state changed, 53
HA UUID, 86
HA:, 34
Half, 228
Hardware error occured, 53
Hardware informantion, 222
Hardware serial number, 108
Heartbeat, 233
High Availability, 9, 85, 89-90, 92, 107, 228-229, 233-234, 240
High availability, 97-98
Host, 135, 140, 145
Host limit, 100
Host:, 32
Hostlist, 127, 131, 152, 154
Hostlists, 122, 124
Hostname, 25, 39, 153, 202-203, 223
Hosts:, 33

I

Idle time before destination is closed, 174
Ignore, 123
Ignore ambiguous program field, 131
Import, 23, 236
Import configuration, 80, 102, 236
Import from file, 124
Include file list, 57
Indexed, 161
Indexed fields, 140
Indexer, 145
Install a new SCB, 242
Installation Steps, 242
Installer, 242
Interface IP, 91-92
Interfaces, 37-39
Interfaces for Heartbeat, 87, 90
Internet Options, 31
Internet Protocol (TCP/IP), 18
Interval, 207
INVALID, 230
Invalidated, 229
IP Address, 20
IP Addresses, 19
IP Settings, 19
IPMI default gateway IP, 108
IPMI IP address, 108
IPMI IP address source, 108
IPMI subnet mask, 108
ISO date, 163

J

Join domain, 153
Join HA, 228-229, 233
Jump to, 199
Jump to last option, 183

K

Key, 59, 114, 177-178
Key >, 192, 194
Key Usage, 119

L

Label, xiii
Last login:, 32
LDAP, 73
Least, 195, 197
Legacy, 130, 160, 163
License, 99-100
License limit reached, 53
License:, 33
Limit of alerts sent out in a batch, 52
Listening address, 129
Load 15:, 34
Load 1:, 34
Load 1|5|15 maximum, 49
Load average, 227
Local, 70, 175
local, 166, 198
Local Area Connection, 17
Local Users, 68, 81
Locality, 27, 113
Locked:, 33
Log, 6, 50-52, 58, 61-62, 82, 103, 127-129, 132, 143, 146, 149, 152, 154, 157, 159, 161, 164, 166-168, 170, 173, 175-177, 194, 204, 211-212, 214, 216, 219, 226, 236
Log Alerts, 223
log-view, 82
log-write, 82
Login failed, 52
Logout, 22, 107, 239
Logout from the management interface, 52
Logs, 35, 103, 140
Logspace, 204
Logspace exceeded warning size, 146, 148
Logspace name, 181
Logspaces, 227
LogStore, 144
Logtype, 224

M

Mail settings, 42, 49, 205-206, 209
Main menu, 31
Make HA IP permanent, 229
Make this extension critical, 119
Manage, 118
Management, 33, 35, 38, 41-42, 44-45, 47, 49, 57-59, 61, 101, 105-106, 110, 112, 114, 122, 146, 148, 177, 205-206, 209, 223-224, 236
Management enabled, 36, 38
Management interface, 37-38, 227
Manual archiving, 204
Master alert, 51
Match, 123
Maximum, 47, 50, 52
Maximum connections, 129
Maximum number of files in notification, 57
Maximum number of search results, 145
Maximum number of statistics to process, 204
MD5 or SHA1, 45-46
Memory, 227
Memory buffer size, 147
Memory limit, 140
Menu, xiii
Message, 83, 140, 145, 204, 213, 216
Message part, 171
Message rate alerting, 50, 135, 137
Message rate alerting statistics, 51-52
Message size, 11
Message throttle, 163
Message:, 224
Messages fetched in a single poll, 5-6
Minimal password strength, 71, 77
Minimum, 50, 52
Modify, 110
Modify User, 110
Modules:, 33
Month, 226
Monthly reports, 206

N

Name, 206, 215
Name resolving, 176
Name/value pairs, 140, 145
Naming, 39
Netmask, 20, 38, 40
Network, 36-37, 39, 246
Network Connections, 17
Network connections, 227
Networks, 38
New, 121
New root password, 106
New value, 83, 202
New Virtual Machine Wizard, 247

Next, 22, 24, 26-27, 117
Next hop IP, 93
Next hop monitoring, 87, 92, 233
NFS, 57, 61, 66, 150-151, 154-155
Nick name, 39
No, 243
No archiving, 60
No encryption, 101
Node HA state, 87
Node HA status, 228
Node HA UUID, 87, 228
Node ID, 87
NOT USED, 230
NTP server, 25
Number of entries, 197
Number of passwords to remember, 77

O

OK, 119, 121, 201, 209, 230
Old value, 83, 202
On, 222
Once, 51-52
Only accept certificates authenticated by the specified CA certificate, 75
Only cleanup, no archiving, 60
Only from persistent configuration, 176
Only with the name, 140
Options, 6, 51-52, 127, 129, 173, 175-177, 204, 211-212, 216, 226, 245
Organization, 27, 113
Organization unit, 28, 113
Other node, 84-85, 91, 93, 97, 233-234
Output disk buffer, 159, 164-165
Output memory buffer, 5, 159, 164-165

P

Page, 83, 202
Password, 65, 69, 153, 158
Password expiration, 70, 77
Password provided by database, 70
Path, 64
Paths, 103, 166-167, 170, 219
Pattern, 216
Pattern Database, 214, 216
Peer configuration, 207
Peer configuration change, 198
Peer Configuration Change, 202
Peer verification, 129
Pending Requests, 116
Per application, 145, 148
Per host, 145, 148
Per host and application, 145, 148
Permanent >, 191
Persistent hostname list, 176
Pid, 140
Pie chart and List, 195
Ping, 153
Ping host, 153
Policies, 55, 59, 82, 122, 124, 150, 152, 236
policies-view, 82
policies-write, 82
Policy, 204
Port, 64, 158, 162, 164, 223
Posix, 74
POSIX group membership attribute name, 75
Preferences, 31, 34
Primary DNS server, 39
Priority, 140
Private keystore, 191, 193
Processes, 227
Production MAC, 90
Program, 135, 140, 145, 203, 213, 216
Program pattern, 215-216
Properties, 17-18
Put all columns into SDATA, 135, 137

Q

Query, 63

R

RADIUS, 76
Raid status, 87
Raid status changed, 53
Raid status:, 33
Rate limit, 127-128
Read old records, 135, 137
Reboot, 96-97, 99, 234, 243
Reboot Cluster, 88, 92-93, 240
Reboot cluster, 89
Recipient, 210
Recommended, 112
Redundant, 90, 228
Redundant Heartbeat status, 90, 228
Remaining time:, 33
Remote, 175
Remote host, 161

Remove, 201
Replace, 124
Replacement value, 171
report, 82
Report from, 207
Report settings, 197
Report subchapter name, 197
Report to, 207
Reporting, 35, 199, 209
Reports, 82, 199, 205, 207-209
Reports are accessible by the following groups, 209
Require commit log, 82-83, 202
Restart syslog-ng, 100, 123, 125
Restore, 149
Restore ALL, 236
Restore now, 236
Retention time, 159
Retention time in days, 60-61
Revert Configuration, 98
Rewrites, 170
RLTP, 129-130
RLTP TLS, 129-130
Root password, 26
Routing table, 38-39
Rsync over SSH, 56, 62
Rule description, 204
Rule ID, 204
Rules, 216
Ruleset name, 213, 216
Run, 116

S

Sampling interval, 204, 226
Save, 78-79
Save As Report subchapter, 197
Save the collected debug info, 225
Scale, 199
Seal the box, 27
Sealed mode, 107
Search, 35, 79-80, 82, 103, 181, 184, 198, 202-204, 213, 216, 221, 223, 227
search, 82
Search in, 227
Secondary DNS server, 39
Security passphrase, 192
Select resolution, 227
Send e-mail alerts to, 43
Send e-mails as, 43
Send even empty reports, 210
Send notification on all events, 57, 61
Send notification on errors only, 57, 61
Send reports in e-mail, 210
Send reports to, 43, 205-206, 209
Sender address, 203
Senders:, 33
Serial, 100
Server Address, 73-75
Server Authentication, 119
Server certificate, 110-111
Server host key, 63
Server private key, 27
Server URL, 175
Server X.509 certificate, 27
Service control, 100, 123, 125, 232
Set, 59, 75, 114
Set Date & Time, 40
Set Default Port, 133
Settings, 17, 70, 73, 76, 82, 202, 247
Severity, 135, 137
SHA-1 fingerprint, 179
Share, 66
Shared secret, 77
Shares, 150, 152
Sharing policy, 146, 148, 152, 154
Shells, 21, 109, 232, 239
Show, 213, 216
Shutdown, 97
Shutdown, Suspend, Reset, 245
Signature, 203
Signature is proof of origin, 119
Size, 149
SMB/CIFS, 56, 60, 62, 64-65
SMB/CIFS options, 150, 152
SMTP server, 25, 42
SMTP server address, 42
SNMP agent settings, 45
SNMP destination, 164
SNMP server address, 44
SNMP settings, 49
SNMP source, 127
SNMP trap settings, 44
SNMP v2c, 44, 164
SNMP v2c agent, 46
SNMP v3, 44, 164
SNMP v3 agent, 46
Source, 167

Sources, 50, 128, 132, 176
Spaces, 50, 58, 61-62, 103, 140, 143, 146, 149, 152, 154, 168, 181, 194, 198, 236
Speed, 37
Split brain, 229-230
Spoof source address, 162
SQL, 132
SSH settings, 105
SSL certificate, 110, 112, 114, 122
SSL certificates, 177
SSL/TLS, 74
Standalone, 228
Standalone mode, 150
Start, 116, 225
Start menu, 17
Start time, 56, 60
STARTTLS, 74
State or Province, 28, 113
Statistics, 197
Status, 86, 228, 233-234
Stop, 225
Submit new request..., 116
Successful login, 52
Suppress timeout, 163
Swap, 202
Swap utilization maximum, 49
Sync Master, 41
Sync now, 41
Sync Slave to Master, 41
Sync source, 229
Sync target, 229
SyncSource, 233
SyncTarget, 233
Syslog, 129-130
Syslog flags, 131
Syslog protocol, 130, 163
Syslog traffic, indexing & search:, 100, 123, 125, 232
syslog-ng, 226
Syslog-ng statistics, 204-205
syslog-ng statistics, 227
syslog-ng traffic statistics, 207
System, 57, 59, 84, 96-102, 107, 123, 125, 222, 232-233, 236
System backup, 58-59, 101, 236
System backup policy, 57-58
System contact, 46
System control, 84, 99
System Control, 96
System debug, 224
System description, 46
System health information, 207
System location, 46
System monitor, 31, 33, 229
System Monitor, 35
System related traps, 50

T

Table, 134
Table of contents, 209
Table rotation, 159
Tags, 140, 195
Tail, 224
Target server, 56, 61
Target settings, 61-62, 65
TCP, 129
Template, 145, 147-148
Template display name, 118
Temporary >, 193
Test, 43, 76
Test connection, 160
Test connection and fetch tables, 133
Test data retrieving, 136-137
Text file, 147
This node, 84-85, 91, 93, 96-97, 99
Time Stamping, 119
Time sync lost, 53
Time-based statistics, 204-205
Time:, 33
Timestamp, 83, 202-204
Timestamp fractions of a second, 159, 164-165
Timestamping error occured, 53
Timestamping frequency, 145
Timezone, 25, 41, 131, 135, 137, 159, 164-165
TLS, 7, 129
TLS certificate, 177
TLS private key, 177-178
TLS settings, 129, 177
Tools, 31
Top, 197
Top/Least statistics, 204
Transport, 129, 162
Troubleshooting, 34, 81, 153, 198, 222-224
Trusted, 129
Trusted distinguished names, 179
Trusted fingerprints, 179
TSA certificate, 110-111

TSA private key, 122
TSA X.509 certificate, 122
Type, 74, 79-80, 144, 147, 151, 154

U

UDP, 129
Unblock Slave Node, 88
Unique ID column, 134
Upload, 24, 59, 75, 100, 102, 114, 122, 124, 144, 177-178, 216
Upload key, 177-178
Use DNS, 128, 131, 176
Use FQDN, 131
use static subchapters, 208
User, 69
User database, 70
User info, 31
User menu, 31-32, 34
User Menu, 191, 193
User:, 32
Username, 44, 46, 63, 65, 153, 158
Username (userid) attribute name, 75
Users, 110

V

Validity, 100, 203
Verify password, 69
Version, 203
Version details, 97, 233
View, 224
View graph, 227
View log files, 223
Visible columns, 185, 201
Visualization, 197
VMware Tools, 245

W

Warning size, 146, 148
Warning, all data on the hard drive(s) will be erased. Are you sure?, 242
Web interface timeout, 33, 35
Web Server, 118
Week, 226
Weekly reports, 206
WFConnection, 230

Y

Year, 226
Yes, 242-243